The Macroeconomic Consequences of Microeconomic Phenomena in the Housing and Labor Markets

The Macroeconomic Consequences of Microeconomic Phenomena in the Housing and Labor Markets

A dissertation presented

by

Adam Michael Guren

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

May 2014

Advisors: Emmanuel Farhi, Edward Glaeser Adam Michael Guren

The Macroeconomic Consequences of Microeconomic Phenomena in the Housing and Labor Markets

Abstract

This dissertation consists of three independent chapters, each of which use microeconomic data and methods to inform an analysis of macroeconomic models and questions. The first two chapters study the short-run dynamics of housing markets, while the last chapter studies fluctuations in labor markets.

The first chapter examines house price momentum, the positive autocorrelation of price changes which in housing markets lasts for two to three years. The chapter introduces, empirically grounds, and quantitatively analyzes an amplification mechanism that can generate substantial momentum from small frictions: sellers with an incentive not to set a unilaterally high or low list price for their house gradually adjust their price so it remains close to the market average. In doing so, the chapter provides new evidence for explanations for price stickiness for which there is little direct evidence. Furthermore, the chapter demonstrates that the resulting momentum helps explain the short-run dynamics of price, volume, and inventory in housing markets.

The second chapter demonstrates how foreclosures can exacerbate a housing bust and delay the housing market's recovery and shows that such effects played a significant role in the recent housing downturn. Foreclosures drive down prices and freeze up the retail (non-foreclosure) market by raising the ratio of sellers to buyers and making buyers more selective. This can push more homeowners underwater and cause more defaults, amplifying an initial shock. When calibrated to the recent housing cycle, the model implies that foreclosures have much larger effects than previously estimated due to general equilibrium spillovers.

The third chapter analyzes why macroeconomic calibrations imply much larger labor supply elasticities than microeconometric studies, paying particular attention to the extensive (participation) margin which is frequently used to explain the divergence. The chapter uses a calibrated macro model to simulate the impacts of tax policy changes on labor supply. It also presents a meta-analysis of quasi-experimental estimates of extensive margin elasticities. Both

approaches show that micro and macro are consistent for steady-state Hicksian elasticities, but micro estimates of extensive-margin Frisch elasticities are an order of magnitude smaller than the values needed to explain business cycle fluctuations in aggregate hours.

Contents

Introduction **1**

1 The Causes and Consequences of House Price Momentum **1**
1.1 Introduction . . . 1
1.2 Four Facts About Housing Dynamics . . . 5
1.2.1 Momentum . . . 5
1.2.2 Housing Cycles . . . 8
1.3 Are Housing Demand Curves Concave? . . . 12
1.3.1 Data . . . 14
1.3.2 Empirical Approach . . . 15
1.3.3 Results . . . 23
1.4 A Model of House Price Momentum . . . 26
1.4.1 Setting . . . 26
1.4.2 The Housing Market . . . 29
1.4.3 Flexible Price Setting . . . 32
1.4.4 Source of Insensitivity 1: Staggered Price Setting . . . 35
1.4.5 Source of Insensitivity 2: A Small Fraction of Rule-of-Thumb Sellers . . . 38
1.5 How Much Can Concave Demand Amplify Momentum? . . . 41
1.5.1 Frequency of Price Adjustment . . . 41
1.5.2 Calibration and Estimation . . . 41
1.5.3 Amplification of Momentum in the Calibrated Model . . . 46
1.6 Can Momentum Help Explain Housing Cycles? . . . 50
1.6.1 Impulse Responses of Non-Price Variables . . . 50
1.6.2 Explaining the Housing Cycle Facts . . . 52
1.7 Conclusion . . . 55

2 How Do Foreclosures Exacerbate Housing Downturns? **57**
2.1 Introduction . . . 57
2.2 Empirical Facts . . . 61
2.3 Housing Market Model . . . 67
2.3.1 Setup . . . 67
2.3.2 Numerical Methods . . . 76
2.4 Basic Model Results and Mechanisms . . . 76
2.4.1 Market Tightness, Choosey Buyer, and Compositional Effects . . . 76
2.4.2 Qualitative Results . . . 78
2.4.3 Isolating the Role of Each Effect . . . 81
2.5 An Extended Model of Default . . . 83
2.5.1 Default in the Extended Model . . . 83
2.5.2 Starting the Downturn . . . 86
2.6 Quantitative Analysis: Calibration and Amplification . . . 87
2.6.1 Calibration . . . 87
2.6.2 Strength of Amplification Channel and Comparative Statics . . . 91
2.7 Cross-MSA Quantitative Analysis . . . 93
2.8 Welfare and Policy Implications . . . 99
2.8.1 Welfare . . . 99
2.8.2 Delaying Foreclosure . . . 102

2.8.3 Interest Rate or Payment Reductions . . . 103
2.8.4 Principal Reduction . . . 104
2.8.5 Other Policies . . . 105
2.9 Conclusion . . . 105

3 Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities 108
3.1 Introduction . . . 108
3.2 Indivisible Labor: Background . . . 113
3.3 Terminology . . . 114
3.4 Simulations of Quasi-Experiments in the RW Model . . . 116
3.5 Meta-Analysis . . . 123
3.6 Comparing Micro and Macro Estimates . . . 127
3.7 Conclusion . . . 133

References 135

A Chapter 1 Appendix 151
A.1 Data . . . 151
A.1.1 Time Series Data . . . 151
A.1.2 Micro Data . . . 154
A.1.3 List Prices Relative to Transaction Prices . . . 161
A.2 Housing Market Facts . . . 163
A.2.1 Momentum . . . 163
A.2.2 Housing Cycle Facts . . . 166
A.2.3 Buyer and Seller Entry . . . 169
A.3 Micro Evidence For Concave Demand . . . 172
A.3.1 Binned Scatter Plots . . . 172
A.3.2 Proof of Lemma 2 . . . 172
A.3.3 Instrumental Variable Robustness and Specification Tests . . . 173
A.3.4 Ordinary Least Squares . . . 182
A.3.5 Monte Carlo Assessment of Bias From Other Sources of Markup Variation . . . 185
A.4 Model . . . 187
A.4.1 Laws of Motion and Value Functions . . . 187
A.4.2 Proofs . . . 190
A.4.3 Frictionless Model . . . 194
A.4.4 Backward-Looking Model . . . 195
A.4.5 Staggered Pricing Model . . . 196
A.4.6 Steady State . . . 198
A.4.7 Simulation Details . . . 199
A.4.8 Understanding Momentum Arising From Staggered Prices . . . 200
A.5 Calibration . . . 202
A.5.1 Calibration Targets . . . 202
A.5.2 Estimation and Calibration Procedure . . . 206
A.5.3 Additional Calibrated Values . . . 209
A.6 Additional Simulation Results . . . 210
A.6.1 Non-Concave Benchmarks . . . 210
A.6.2 Downward Shock . . . 210

A.6.3 Impulse Responses For All Variables . . . 211
A.6.4 Calibration Robustness: Lower Seller Search Cost . . . 212
A.6.5 Deterministic Impulse Responses For Staggered Model . . . 213

B Chapter 2 Appendix **216**
B.1 Derivations and Proofs . . . 216
B.1.1 Steady Sate . . . 216
B.1.2 Dynamics . . . 219
B.1.3 Calibration . . . 221
B.1.4 Extended Model . . . 222
B.2 Details Omitted From Main Text . . . 224
B.2.1 Details of Isolating Each Effect . . . 224
B.2.2 Cross-Markets Analysis With 10% REO Discount . . . 224
B.3 Judicial vs. Non-Judicial States . . . 226
B.4 Data Sources and Calculations . . . 227

C Chapter 3 Appendix **232**
C.1 Simulations of Quasi-Experiments (Figure 24) . . . 232
C.2 Meta-Analysis of Quasi-Experimental Estimates (Table 17) . . . 235
C.3 Micro vs. Macro Elasticities (Table 18) . . . 238
C.4 Technical Appendix . . . 239

Acknowledgements

I would like to thank my advisors, professors John Campbell, Raj Chetty, Emmanuel Farhi, Edward Glaeser, Nathan Hendren, and Lawrence Katz for their patience, advice, and guidance over the past six years. Their influence can be seen throughout this dissertation, and the lessons they have taught me will continue to shape my work for years to come.

I also benefited from a number of faculty both at Harvard and visiting from other institutions who had an important impact on this dissertation and my way of thinking. These include Joe Altonji, Pat Bayer, Gary Chamberlain, John Friedman, James Hamilton, Sam Hanson, Erik Hurst, Max Kasy, Greg Mankiw, Chris Mayer, Jonathan Parker, Alp Simsek, Andrei Shleifer, and Lawrence Summers. I also received valuable feedback from many other participants at seminars and workshops. Erik Hurst in particular deserves a special acknowledgement for the role he played in leading me towards empirical macroeconomics.

I would also like to thank my co-authors. The second chapter is joint work with Tim McQuade and is truly a team effort. The third chapter is joint work with Raj Chetty, Day Manoli, and Andrea Weber and has been published in the 2012 NBER Macroeconomics Annual.

I received extensive assistance in understanding the data used in this dissertation from Andrew Cramond, Kathryn Dobbyn, T.J. Doyle, Mark Fleming, Sam Khater, and Mike Simonsen. Shelby Lin and Christopher Palmer assisted me in formatting and cleaning data.

Perhaps the most formative part of my graduate school experience has been countless discussions with my peers Nikhil Agarwal, Rebecca Diamond, Michal Kolesar, and Tim McQuade, not only in classrooms and in front of white boards but also at meals, over coffee, and in our offices. They made graduate school fun and will be my lifelong friends. I was also influenced by many other peers, a few of whom deserve special mention: Eduardo Davila, Will Diamond, Will Dobbie, Peter Ganong, David Hemous, Ben Hebert, Yuhta Ishii, Brock Mendel, Morten Olsen, Pascal Noel, Mikkel Plagborg-Moller, Ben Schoefer, and Danny Yagan.

Finally, I would like to thank Kate Lewandowski for her love, support, and patience through the process of writing this dissertation and my parents Marc and Aliza and sister Julia for always being there for me through good times and bad and for teaching me most of what I know.

Introduction

This dissertation consists of three independent chapters, each of which use microeconomic data and methods to inform an analysis of macroeconomic models and questions. The first two chapters study the short-run dynamics of housing markets, while the last chapter studies fluctuations in labor markets.

The first chapter studies house price momentum, the positive autocorrelation of price changes which in the housing market lasts for two to three years. The chapter introduces, empirically grounds, and quantitatively analyzes an amplification mechanism that can generate substantial momentum from small frictions and demonstrates that the resulting momentum helps explain the short-run dynamics of housing markets. The amplification is due to a concave demand curve in relative price, which implies that increasing the quality-adjusted list price of a house priced above the market average rapidly reduces its probability of sale, but cutting the price of a below-average priced home only slightly improves its chance of selling. This creates a strategic complementarity that incentivizes sellers to set their list price close to others'. Consequently, frictions that cause slight insensitivities to changes in fundamentals lead to prolonged adjustments because sellers gradually adjust their price to stay near the average. I provide new micro empirical evidence for the concavity of demand—which is often used in macro models with strategic complementarities—by instrumenting a house's relative list price with a proxy for the seller's equity. I find significant concavity, which I embed in an equilibrium housing search model in which buyers avoid visiting houses that appear overpriced. I demonstrate and quantitatively evaluate the model's ability to amplify two frictions: staggered pricing and a fraction of backwards-looking rule-of-thumb sellers. Both frictions are amplified substantially, and the model explains the momentum observed empirically with a small fraction of rule-of-thumb sellers. Strong house price momentum leads households to re-time their purchase or sale, thereby explaining several features of the dynamic relationships between price, volume, inventory, and buyer and seller entry.

The second chapter, which is joint work with Tim McQuade analyzes the role foreclosures play in housing downturns. The recent housing bust precipitated a wave of mortgage defaults, with over seven percent of the owner-occupied housing stock experiencing a foreclosure. This chapter presents a model that shows how foreclosures can exacerbate a housing bust and delay the housing market's

recovery. By raising the ratio of sellers to buyers, by making buyers more selective, and by changing the composition of houses that sell, foreclosures freeze up the market for retail (non-foreclosure) sales and reduce both price and volume. Because negative equity is necessary for default, these general equilibrium effects on prices can create price-default spirals that amplify an initial shock. To assess the magnitude of these channels, the model is calibrated to simulate the downturn. The amplification channel is significant. The model successfully explains aggregate and retail price declines, the foreclosure share of volume, and the number of foreclosures both nationwide and across MSAs. While the model can explain variation in sales across MSAs, it cannot account for the aggregate level of the volume decline, suggesting that other forces have reduced sales nationwide. The quantitative analysis implies that from 2007 to 2011 foreclosures exacerbated aggregate price declines by approximately 50 percent and declines in the prices of retail homes by approximately 30 percent.

The third chapter, which is joint work with Raj Chetty, Day Manoli, and Andrea Weber, assesses why macroeconomic calibrations imply much larger labor supply elasticities than microeconometric studies. One prominent explanation for this divergence is that indivisible labor generates extensive margin responses that are not captured in micro studies of hours choices. We evaluate whether existing calibrations of macro models are consistent with micro evidence on extensive margin responses using two approaches. First, we use a standard calibrated macro model to simulate the impacts of tax policy changes on labor supply. Second, we present a meta-analysis of quasi-experimental estimates of extensive margin elasticities. We find that micro estimates are consistent with macro evidence on the steady-state (Hicksian) elasticities relevant for cross-country comparisons. However, micro estimates of extensive-margin elasticities are an order of magnitude smaller than the values needed to explain business cycle fluctuations in aggregate hours. Hence, indivisible labor supply does not explain the large gap between micro and macro estimates of intertemporal substitution (Frisch) elasticities. Our synthesis of the micro evidence points to Hicksian elasticities of 0.3 on the intensive and 0.25 on the extensive margin and Frisch elasticities of 0.5 on the intensive and 0.25 on the extensive margin.

Chapter 1:
The Causes and Consequences of House Price Momentum

1.1 Introduction

A puzzling and prominent feature of housing markets is that aggregate price changes are highly positively autocorrelated, with a one percent annual price change correlated with a 0.30 to 0.75 percent change in the subsequent year (Case and Shiller, 1989).[1] This price momentum lasts for two to three years before prices mean revert, a time horizon far greater than most other asset markets. Substantial momentum is surprising because predictable price changes should be arbitraged away by investors and households that can re-time their purchase or sale and because most pricing frictions dissipate quickly.

This chapter introduces, empirically grounds, and quantitatively analyzes an amplification mechanism that can generate substantial momentum from small frictions. The mechanism relies on a strategic complementarity among list-price-setting sellers that makes the optimal list price for a house depend positively on the prices set by others (Cooper and John, 1988). Strategic complementarities of this sort are frequently used in macroeconomic models (*e.g.*, Ball and Romer, 1990; Woodford, 2003; Angeletos and La'O, 2013) but there is limited empirical evidence of their importance and strength. In analyzing momentum in the housing market, I provide micro empirical evidence for a prevalent strategic complementarity in the macroeconomics literature and, using a calibrated equilibrium search model, demonstrate that its ability to amplify underlying frictions is quantitatively significant.

I also show that momentum has important consequences that help explain several perplexing features of the dynamics of housing markets relating to sales and inventory in addition to price. These dynamics, which are analogous to several features of business cycles, matter for the macroeconomy because housing markets affect household balance sheets, the financial system, and business cycles and are a potential channel for monetary policy. House price momentum may also explain why recoveries from housing-triggered cycles are slow.

The propagation mechanism I introduce relies on two components: costly search and a demand

[1] See also Cutler et al. (1991), Abraham and Hendershott (1996), Cho (1996), Malpezzi (1999), Meen (2002), Capozza et al. (2004), Head et al. (2014), and Glaeser et al. (2013).

curve that is concave in relative price. Search is inherent to housing because no two houses are alike and idiosyncratic taste can only be learned through costly inspection. Search and idiosyncratic taste also limit arbitrage by creating endogenous transaction costs and by making the market price for a house difficult to ascertain. Concave demand in relative price implies that the probability a house sells is more sensitive to list price for houses priced above the market average than below the market average. While concave demand may arise in housing markets for several reasons, I focus on the manner in which asking prices direct buyer search. The intuition is summarized by an advice column for sellers: "Put yourself in the shoes of buyers who are scanning the real estate ads...trying to decide which houses to visit in person. If your house is overpriced, that will be an immediate turnoff. The buyer will probably clue in pretty quickly to the fact that other houses look like better bargains and move on."[2] In other words, the probability that a house is visited by buyers decreases rapidly as a home's list price rises relative to the market average. This generates a concave demand curve in relative price because at high relative prices buyers are on the margin of looking and purchasing, while at low relative prices they are only on the margin of purchasing.

Concave demand incentivizes list-price-setting sellers—who have market power due to search frictions—to set their list prices close to the mean. Intuitively, raising a house's relative list price reduces the probability of sale and profit dramatically, while lowering its relative price increases the probability of sale slightly and leaves money on the table. Modest frictions that generate initial insensitivities of prices to changes in fundamentals cause protracted price adjustments because sellers find it optimal to gradually adjust their price so that they do not stray too far from the market average.

To evaluate the concavity of the effect of unilaterally changing a house's relative quality-adjusted price on its sales probability, I turn to micro data on listings for the San Francisco Bay, Los Angeles, and San Diego metropolitan areas from 2008 to 2013. I address bias caused by unobserved quality by instrumenting relative list price with the amount of aggregate price appreciation since the seller purchased. The identification strategy takes advantage of the fact that sellers with low appreciation since purchase set higher list prices because the equity they extract from the sale of their current home constrains their ability to make a down payment on their next home (Stein, 1995; Genesove

[2] "Settling On The Right List Price for Your House," Ilona Bray, http://www.nolo.com/legal-encyclopedia/listing-house-what-price-should-set-32336-2.html.

and Mayer, 1997). Because I compare listings within a ZIP code and quarter, this supply-side variation identifies the curvature of demand if unobserved quality is independent of when a seller purchased their home. The instrumental variable estimates reveal a concave relationship that is statistically and economically significant.[3] My findings about the concavity of demand are robust to other sources of relative price variation that are independent of appreciation since purchase.

To assess the strength of this propagation mechanism, I embed concave demand in a Diamond-Mortensen-Pissarides equilibrium search model. I explore the effects of two separate sources of price insensitivity. First, I consider staggered pricing whereby overlapping groups of sellers set prices that are fixed for multiple periods (Taylor, 1980). Concave demand induces sellers to only partially adjust their prices when they have the opportunity to do so, and repeated partial adjustment manifests itself as additional momentum. Second, I introduce a small fraction of backward-looking rule-of-thumb sellers as in Campbell and Mankiw (1989) and Gali and Gertler (1999). Backward-looking expectations are frequently discussed as a potential cause of momentum (*e.g.*, Case and Shiller, 1987; Case et al. 2012), but some observers have voiced skepticism about widespread non-rationality in housing markets given the financial importance of housing transactions for most households. With a strategic complementarity, far fewer backward-looking sellers are needed to explain momentum because the majority of forward-looking sellers adjust their prices gradually so they do not deviate too much from the backward-looking sellers (Haltiwanger and Waldman, 1989; Fehr and Tyran, 2005). This, in turn, causes the backward-looking sellers to observe more gradual price growth and change their price by less, creating a two-way feedback that amplifies momentum.

I calibrate the parameters of the model that control the shape of the demand curve to match the micro empirical estimates and the remainder of the model to match steady state and time series moments. The calibrated model generates substantial amplification of the underlying frictions. With staggered pricing, the model can explain a ten month price adjustment—or about one quarter of the momentum in the data—in response to a shock to fundamentals even though all sellers have reset their price within two months of the shock. With rule-of-thumb sellers, the model generates three years of positively autocorrelated price changes as observed empirically if 26.5 percent of sellers are backward-looking. By contrast, without concave demand, 78 to 93 percent of sellers

[3] Although endogeneity is a worry, the ordinary least squares relationship is also concave. However, as one would expect if unobservable quality is an issue, it has a smaller slope.

would have to be backward-looking to generate a three-year response.

The amplification mechanism adapts two ideas from the macro literature on goods price stickiness to frictional asset search. First, the concave demand curve is similar to "kinked" demand curves (Stiglitz, 1979; Woglom, 1982) which, since the pioneering work of Ball and Romer (1990) has been frequently cited as a potential source of real rigidities. In particular, a "smoothed-out kink" extension of Dixit-Stiglitz preferences proposed by Kimball (1995) is frequently used to tractably introduce real rigidities through strategic complementarity in price setting. Second, the repeated partial price adjustment caused by the strategic complementarity is akin to Taylor's (1980) "contract multiplier." A lively literature has debated the importance of strategic complementarities and kinked demand in particular for propagating goods price stickiness by analyzing calibrated models (*e.g.,* Chari et al., 2000), by assessing whether the ramifications of strategic complementarities are borne out in micro data (Klenow and Willis, 2006; Bils et al., 2012), and by examining exchange-rate pass through for imported goods (*e.g.,* Gopinath and Itshoki, 2010; Nakamura and Zerom, 2010). My analysis of housing markets adds to this literature by directly estimating a concave demand curve and assessing its ability to amplify frictions in a calibrated model.

Having established a propagation mechanism for house price momentum empirically and theoretically, I show that momentum affects the dynamics of sales volume and the inventory of houses for sale. Forward-looking buyers and sellers re-time their purchase decisions due to expectations of predictable future price changes. Such re-timing causes sudden swings in inventory that drive the reversal between a hot market, with a substantial excess of buyers, and a cold market, with a relative dearth of buyers. For instance, at a trough, marginal buyers rush to purchase before prices rise, while marginal sellers wait to obtain a better price for their home, leading inventory to plummet.[4] To formalize this story, I build on Novy-Marx (2009) by including buyer and seller entry decisions in the model.

Forward-looking entry responses in the calibrated model help explain three puzzling features of housing cycles. First, seller entry remains high as volume plummets at peaks and remains low as

[4] Buyer and seller quotes in newspapers provide suggestive evidence of such re-timing. In 2013, when prices were rising, a buyer explained to the *Wall Street Journal* "if you don't get in now, things are going to skyrocket over the next year," while a seller who delayed putting their house on the market told the *Journal* that "the extra money – that was worth [waiting] for the year." This effect is part of the folk wisdom of housing markets, yet has not appeared in the academic literature. For instance, Calculated Risk Blog describes a conversation with a real estate agent who argues that "In a market with falling prices, sellers rush to list their homes, and inventory increases. But if sellers think prices have bottomed, then they believe they can be patient, and inventory declines."

volume picks up at troughs, which is the exact pattern created by the re-timing of entry in light of momentum. Second, volume and inventory are more volatile than price. This is difficult to reconcile with most calibrations of housing search models in a direct analogue to Shimer's (2005) unemployment volatility puzzle for labor search models. With momentum, volume and inventory are more volatile not only because price responds gradually but also because the adjustment of inventory is accelerated by the re-timing of entry. Third, in the data, price changes are strongly negatively correlated with inventory levels (Peach, 1983). This "housing Phillips curve" is surprising because in most asset pricing models, price changes are correlated with changes in fundamentals such as inventory (Caplin and Leahy, 2011). In my model, the quick response of inventory and gradual response of price create a strong correlation between price changes and inventory levels.

The remainder of the chapter proceeds as follows. Section 1.2 introduces facts about housing dynamics. Section 1.3 analyzes micro data to assess whether housing demand curves are concave. Section 1.4 presents the model. Section 1.5 calibrates the model to the micro estimates and assesses the degree to which strategic complementarities amplify momentum. Section 1.6 discusses the consequences of this momentum for housing cycles. Section 1.7 concludes.

1.2 Four Facts About Housing Dynamics

1.2.1 Momentum

Since the pioneering work of Case and Shiller (1989), price momentum has been considered one of the most puzzling features of housing markets. While other financial markets exhibit momentum, the housing market is unusual for the strength of the effect and the horizon over which it persists.[5]

Fact 1: *Price changes are serially correlated for 8 to 14 quarterly lags.*

House price momentum has consistently been found across cities and countries, time periods, and price index measurement methodologies (Cho, 1996). Figure 1 shows three measures of momentum

[5] Note that the "momentum" I analyze refers to autocorrelation in aggregate price time series, which is distinct from the short-term over-performance of stocks that recently performed best that is also called "momentum." Time-series momentum holds for a number of other asset classes over shorter horizons. Cutler et al. (1991) look across a large number of asset classes and find that for the vast majority of assets, positive autocorrelation in returns lasts for less than a year. Moskowitz et al. (2012) find that time series momentum lasts for approximately 12 months for 58 different equity index, currency, commodity, and bond futures. This 12 month horizon is an upper bound for the type of momentum studied here, which includes only capital gains, because the measured returns in Moskowitz et al. include both dividends (which are known to be autocorrelated) and capital gains.

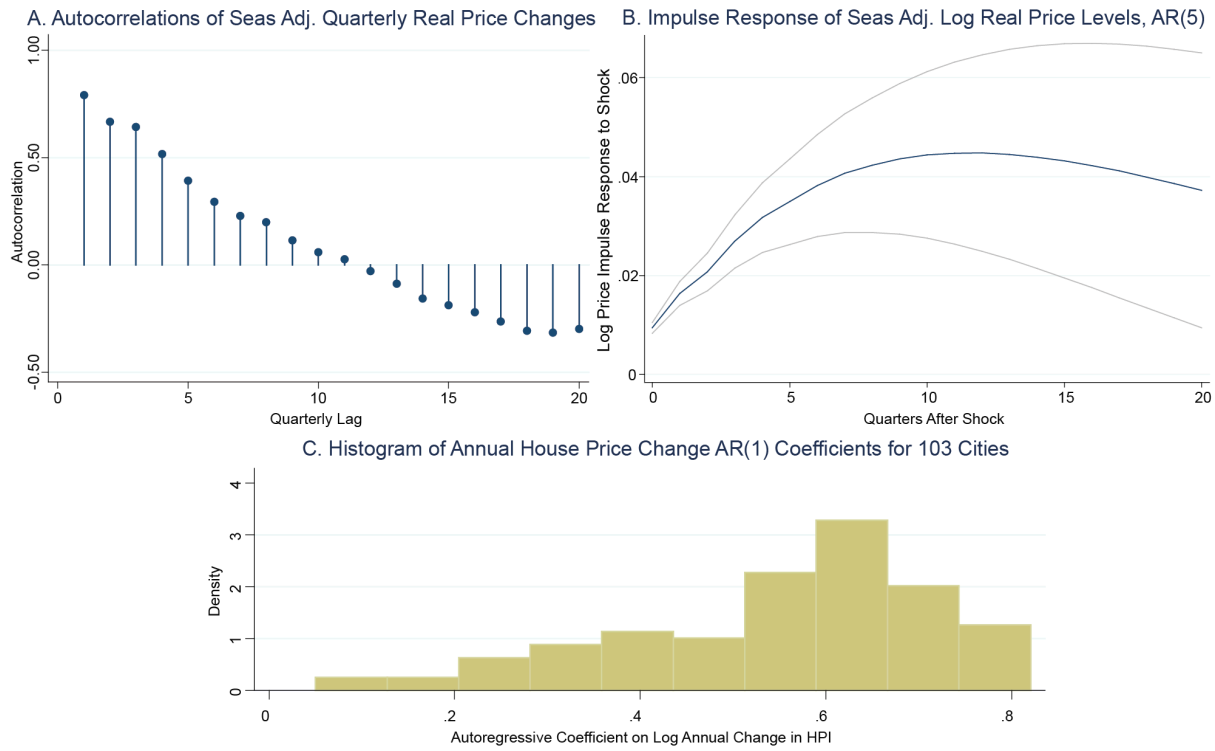

Figure 1: Momentum in Housing Prices

Notes: Panel A and B show the autocorrelation function for quarterly real price changes and an impulse response of log real price levels estimated from an AR(5) model, respectively. The IRF has 95% confidence intervals shown in grey. An AR(5) was chosen using a number of lag selection criteria, and the results are robust to altering the number of lags. Both are estimated using the CoreLogic national repeat-sales house price index from 1976-2013 collapsed to a quarterly level, adjusted for inflation using the CPI, and seasonally adjusted. Panel C shows a histogram of annual AR(1) coefficients of annual house price changes as in regression (1) estimated separately on 103 CBSA division repeat-sales house price indices provided by CoreLogic. The local HPIs are adjusted for inflation using the CPI. The 103 CBSAs and their time coverage, which ranges from 1976-2013 to 1995-2013, are listed in Appendix A.1.

for the CoreLogic national repeat-sales house price index for 1976 to 2013.[6] Panel A shows that autocorrelations are positive for 11 quarterly lags of the quarterly change in the price index adjusted for inflation and seasonality. Panel B shows an impulse response in log levels to an initial one percent price shock estimated from an AR(5). In response to the shock, prices gradually rise for two to three years before mean reverting. Finally, panel C shows a histogram of AR(1) coefficients estimated separately for 103 metropolitan area repeat-sales house price indices from CoreLogic using a regression of the annual change in log price on a one-year lag of itself as in Case and Shiller

[6] As discussed in Appendix A.2, price indices that measure the median price of transacted homes display momentum over roughly two years as opposed to three years for repeat-sales indices.

(1989):

$$\Delta_{t,t-4} \ln p = \beta_0 + \beta_1 \Delta_{t-4,t-8} \ln p + \varepsilon. \quad (1)$$

β_1 is positive for all 103 cities, strongest for cities with inelastic housing supply, and the median city has an annual AR1 coefficient of 0.60. Appendix A.2 replicates these facts for a number of countries, price series, and measures of autocorrelation and consistently finds two to three years of momentum.[7]

The existing evidence suggests that momentum cannot be explained by serially correlated changes in fundamentals. Case and Shiller (1989) argue that momentum cannot be explained by autocorrelation in interest rates, rents, or taxes. Glaeser et al. (2013) estimate a dynamic spatial equilibrium model and find that "there is no reasonable parameter set" consistent with short-run momentum. Capozza et al. (2004) find significant momentum after accounting for six comprehensive measures of fundamentals in a vector error correction model.

Four main explanations have been offered for momentum in asset markets and for the housing market more specifically. First, a behavioral finance literature hypothesizes that investors initially underreact to news due to behavioral biases (Barberis et al., 1998, Hong and Stein, 1999) or loss aversion (Frazzini, 2006) and then "chase returns" due to extrapolative expectations about price appreciation.[8] Both extrapolative expectations and loss aversion are considered to be important forces in the housing market (Case and Shiller, 1987; Berkovec and Goodman, 1996; Glaeser et al., 2013; Genesove and Mayer, 2001). Second, Anenberg (2013) shows that gradual learning about market conditions by sellers can create momentum. Third, Head et al. (2014) demonstrate that strong search frictions and a gradual construction response can cause the liquidity of houses to adjust slowly in response to a shock to local incomes, which creates momentum.[9] Finally,

[7] In the housing market, the price level appears to be sticky but the rate of change does not appear to react sluggishly. In particular, neither the evidence presented here nor the structural panel VAR in Head et al. (2014) shows evidence of autocorrelations of house price changes near one or delayed "hump shaped" impulse responses of house price changes. This is unlike the CPI or GDP deflator, which demonstrate considerable persistence in the rate of change (Fuhrer, 2011).

[8] Frazzini argues that as prices rise, potential sellers who resist selling at a loss relative to their initial purchase begin to experience gains. This causes them to sell, putting downward pressure on prices. A similar point could be made with respect to underwater homeowners who regain positive equity as prices rise.

[9] Head et al. (2014) assume that searching buyers need housing, which must be built when a metropolitan area grows due to an income shock. In their calibrated model, market tightness takes nearly six years to adjust to adjust to a shock due to a slow construction response, search frictions, and a shock that exhibits persistent changes. The gradual adjustment of market tightness creates momentum. By contrast, in the calibrated search model without additional frictions presented here, market tightness adjusts in two years and creates a tiny amount of momentum.

momentum could result from a gradual spread of optimism if sentiment drives house prices rather than fundamentals (Burnside et al., 2013).

Learning and search have been calibrated quantitatively and cannot explain the full extent of momentum in housing markets. Anenberg's (2013) structural model of learning can explain an annual AR(1) coefficient of 0.124 relative to between 0.3 and 0.75 in the data. Head et al.'s (2014) calibrated model can explain half of the autocorrelation in prices at one year, but almost none at two years. In Section 1.5, I show that the vast majority of sellers would have to have a simple form of extrapolative expectations to fully explain the amount of momentum in the data. The role of sentiment has yet to be measured. The amplification mechanism I propose complements these existing explanations by strengthening them so they can better fit the momentum that is empirically observed.

1.2.2 Housing Cycles

I relate three other facts and puzzles about the short-run dynamics of housing cycles to momentum.

Fact 2: *Seller entry rises above sales volume at peaks and falls below sales volume at troughs, corresponding to large and sudden fluctuations in inventory.*

Although sales and seller entry track one another, Figure 2 shows that at the peak of the recent boom and bust cycle, seller entry remained high for several quarters as volume began to plunge, which corresponded to a sudden increase in inventory. Conversely, as volume and prices began to rise in 2012 and 2013, seller entry remained low, coinciding with a sudden drop in inventory. Appendix A.2 shows that this fact is not unique to 2003 to 2013. Although there is no data on the stock of buyers, most models imply that if seller entry lags sales, buyer entry must lead sales. Figure 2 illustrates this by using a simple matching function parameterized based on Genesove and Han (2012) to infer the stock of buyers from sales volume and the stock of homes for sale.

Fact 3: *At an annual frequency, the volatility of sales volume is twice that of real price and the volatility of inventory as measured by months of supply is three times that of real price.*

Despite the predictability of price changes, the housing market is volatile. Table 1 shows the standard deviation of annual log changes for four series: real disposable personal income, real house prices, sales volume, and "for sale" inventory measured as months of supply (a common metric in the housing market). Price is four times more volatile than income, and volume and inventory

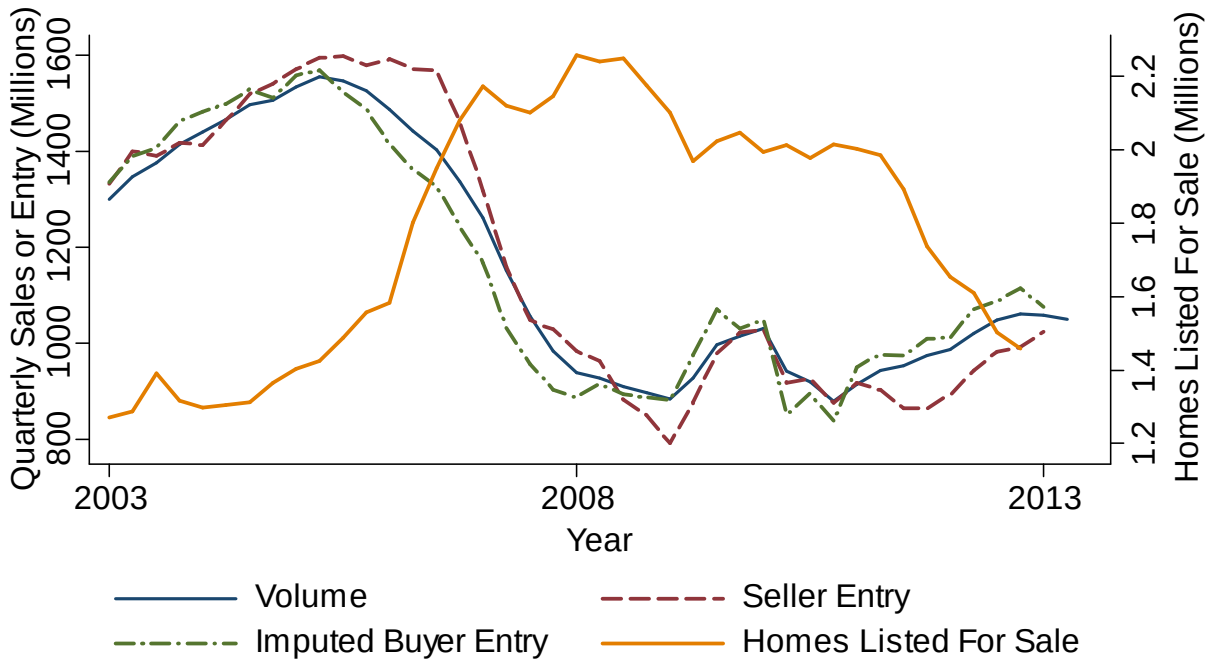

Figure 2: Sales, Entry and Inventory, 2003-2013

Notes: Volume is raw data from the National Association of Realtors of sales of existing single-family homes at a seasonally-adjusted annual rate. Homes listed for sale is from the Census Vacancy Survey. Seller entry is computed as $\text{Entrants}_t = \text{Sellers}_t - \text{Sellers}_{t-1} + \text{Sales}_t$. Buyer entry is computed similarly, but since there is not a raw data series for the stock of buyers it is imputed using a simple Cobb-Douglas matching function $\frac{Sales}{S} = \xi \left(\frac{B}{S}\right)^{-.8}$ with the 0.8 elasticity from Genesove and Han (2012). In this figure, $\xi = 1$ so that in a steady state there is 3 months of supply. All four series are smoothed using a three-quarter moving average.

are, in turn, more volatile than price. The volatility of inventory in particular is of note because substantial fluctuations in inventory at peaks and troughs herald rapid changes between buyers' and sellers' markets. Finally, price and volume are highly positively correlated and both are positively correlated with income.[10]

Fact 4: *(Housing Phillips Curve) Price changes are negatively correlated with inventory levels, with a one log point increase in months of supply correlated with a 0.14 log point decrease in annual price growth (Peach, 1983; Lazear, 2010; Caplin and Leahy, 2011).*

Figure 3 shows the time series of the annual change in the log price index and of log inventory,

[10] There is a substantial literature on the positive correlation of price and sales volume. Stein (1995) and Ortalo-Magne and Rady (2006) argue prices affect the ability of homeowners to extract equity and make a down payment on their next home, leading to a feedback from prices to volume. A literature initiated by Wheaton (1990) and Krainer (2001) uses a steady state Diamond-Mortensen-Pissarides search model to argue that the relative number of buyers and sellers in the market affects the liquidity of homes, creating a feedback from volume to price. Leamer (2007) and Case (2008), among others, suggest that volume is more volatile than price because nominal loss aversion and backwards-looking expectations on the part of sellers make prices sticky downward.

Table 1: Cyclical Summary Statistics for Income, House Price, Sales, and Inventory

	$\sigma_{\log x_q - \log x_{q-4}}$	ρ_x, Real HPI	ρ_x, Sales Volume	ρ_x, Months of Supply
Real Disposable Pers. Income	0.016	.819	.668	.497
Real House Price Index	0.065		.726	.305
Sales Volume	0.143			-.263
Inventory: Months of Supply	0.207			

Notes: All series are for 1976-2013 at a quarterly frequency. The first column shows the standard deviation of annual log changes, while the other columns show correlation coefficients of log levels at a quarterly frequency. Real disposable personal income is BEA series DPIC96. Real price is the CoreLogic national repeat-sales house price index adjusted for inflation using the CPI. Sales volume is from the National Association of Realtors single-family existing home series. Months of supply is created by dividing homes listed for sale from the Census Vacancy Survey by the NAR sales series. The volume, income, and months of supply series are all seasonally adjusted.

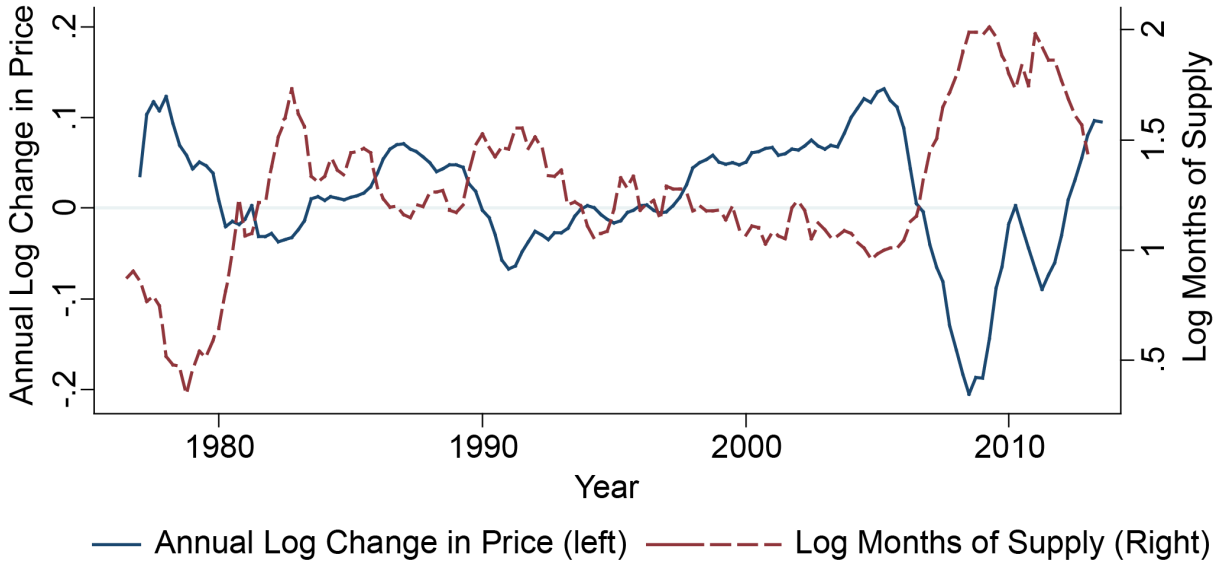

Figure 3: Price Changes Correlated With Inventory Levels

Notes: The figure shows the time series of the annual change in the log CoreLogic national repeat-sales house price index plotted against log months of supply. The latter is calculated by dividing homes vacant for sale from the Census Vacancy Survey by sales of existing single-family homes from the National Association of Realtors (NAR), measured at the midpoint of the yearlong period over which the change in price is computed.

measured as months of supply at the midpoint of the year over which the change in price is calculated. This relationship is reminiscent of the Phillips curve as it relates inventory—the equivalent of unemployment in the housing market—to price appreciation. The visible inverse co-movement in the series is confirmed by a regression: a one percent increase in months of supply is associated with a 0.14 percent decrease in the annual change in prices with an R-squared of 0.53.[11] This

[11] Both the numerator of months of supply—homes for sale—and the denominator—volume—matter. Appendix

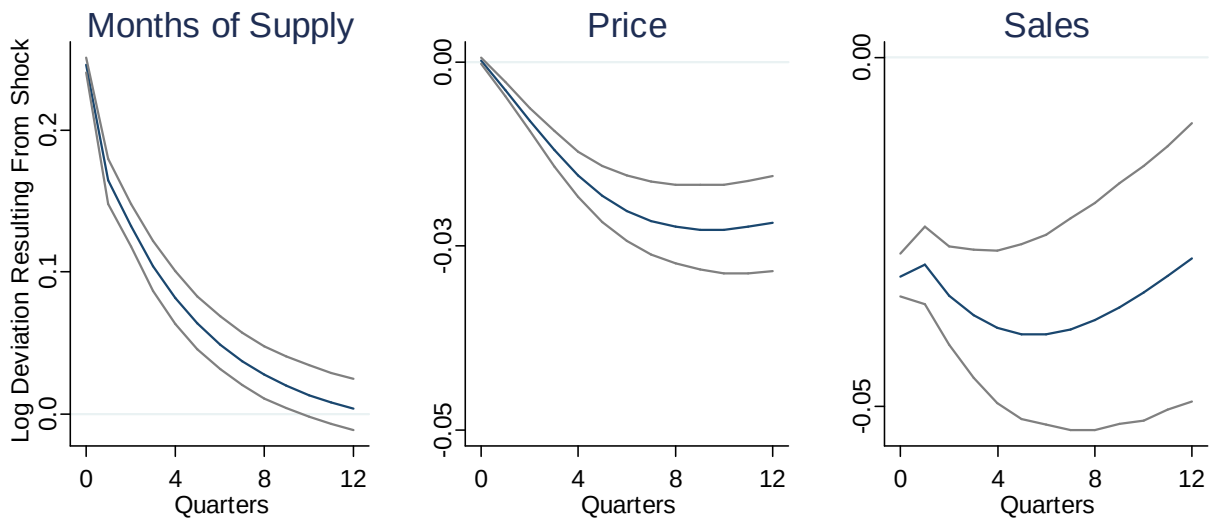

Figure 4: Impulse Response to Inventory Shock in Panel VAR

Notes: The figure shows orthogonalized impulse response functions to a months of supply shock computed from a two-lag panel vector autoregression of log months of supply, log price, and log sales volume for a panel of 42 cities from 1990 to 2013 described in Appendix A.1. Price and sales are from CoreLogic, with price corresponding to the local CoreLogic house price index adjusted for the CPI and sales corresponding to existing home sales. Months of supply at the MSA level comes from the National Association of Realtors. All data is seasonally adjusted, and the panel VAR, which includes a fixed effect for each city as described in Appendix A.2, is estimated using system GMM and Helmert mean differencing using a Stata package by Inessa Love. The OIRFs are computed using a Cholesky decomposition with the variables ordered so that months of supply is assumed not to depend contemporaneously on shocks to price or volume and price is assumed not to depend contemporaneously on shocks to volume. The results are robust as long as months of supply is prior to volume in the Cholesky ordering. The blue line is the OIRF, and the grey bands indicate 95 percent confidence intervals computed using a Monte Carlo procedure that generates 500 impulse responses from draws from the distribution of coefficients implied by the estimated coefficients and their variance-covariance matrix.

relationship is puzzling because most asset pricing models imply price changes should be correlated with changes in variables that reflect fundamentals, such as inventory, rather than with their levels. With mean reverting shocks, such models imply a positive correlation between price changes and inventory levels because when inventory levels are high, inventories tend to fall and prices tend to rise. Caplin and Leahy (2011) show that this effect can be eliminated if prices are posted before shocks are realized.

To bring together the facts, I estimate a panel vector autoregression with city fixed effects on log price, log volume, and log inventory using a panel of 42 cities from 1990 to 2013 described in Appendix A.1. The panel vector autoregression (VAR) is estimated using system GMM as

A.2 shows the time series look similar if log homes listed for sale adjusted for a linear time trend replaces months of supply and that a regression of the annual change in log price on log homes listed for sale has an R-squared of 0.40. By contrast, regressing changes on changes gives a weak correlation.

in Arellano and Bover (1995). Figure 4 shows the orthogonalized impulse response functions of months of supply, price, and sales in response to a one standard deviation positive shock to months of supply. This can be thought of as a negative demand shock because the positive co-movement of price and volume implies that demand-side shocks are predominant. Volume immediately falls, and months of supply monotonically and gradually decays back to its steady state value after the initial shock. Prices gradually decline over a 10-quarter period, with the price decline tapering off as inventory returns to steady state. Appendix A.2 shows similar results for a VAR and a vector error correction model estimated on national data.

The model presented in this chapter implies that housing cycles with these four features arise from the interaction of small underlying frictions, strategic complementarities, and forward-looking decisions about when in the cycle to buy and sell. Several other papers have discussed how the endogenous timing of purchasing decisions affects housing cycles, although not in light of momentum. Novy-Marx (2009) shows that entry responses can amplify the long-run response to shocks and increase the amplitude of cycles. Anenberg and Bayer (2013) demonstrate that the cost of simultaneously holding two homes in an illiquid market can make the number of households who simultaneously buy and sell pro-cyclical, which increases volatility.

More directly related to this chapter are explanations for why seller entry may fall at troughs and rise at peaks. One reason why seller entry may remain low as volume rises after a trough is nominal loss aversion (Genesove and Mayer, 2001) and lock in due to negative equity (Stein, 1995). Head et al. (2014) present another mechanism: when local incomes rise, new entrants to an MSA need a place to live, which drives up rents until new homes are built and causes potential sellers to rent their homes temporarily before selling then. More broadly, Head et al. (2014) is most closely related to this research. Their analysis of the joint responses of construction, house prices, house sales, and population to city-level income shocks in a model with momentum is complementary to my focus on the timing of purchase and sale decisions of existing homeowners and residents.

1.3 Are Housing Demand Curves Concave?

I propose an amplification channel for momentum based on search and a concave demand curve in relative price. Search is a natural assumption for housing markets, but the relevance of concave demand requires further explanation.

A literature in macroeconomics shows how strategic complementarities among goods producers can amplify small pricing frictions into substantial price sluggishness by incentivizing firms to set prices close to one another. Strategic complementarities operate either through a monopolistic firm's marginal cost or its markup, which pushes a firm to price close to the market average if demand is concave in relative price. "Kinked demand" was introduced by Stiglitz (1979) and Woglom (1982), who hypothesized that firms that increase their price induce consumers to search for a new firm, but firms that cut their price only gain a few active searchers. Ball and Romer (1990) show that this can create real rigidities and possibly explain why prices are so sticky despite small menu costs. This argument has been formalized in several papers, such as Benabou (1992) and Levin and Yun (2009). Kimball (1995) generalizes Dixit-Stiglitz-style aggregator to allow for concave demand, which is used as an important real rigidity in several popular New Keynesian models (*e.g.*, Smets and Wouters, 2007). Despite the frequency with which it is used, there is little direct evidence for concave demand.[12]

Because momentum is similar to price stickiness in goods markets, I hypothesize that a similar strategic complementarity may amplify house price momentum. There are several reasons why concave demand may arise in housing markets. First, buyers may avoid visiting homes that appear to be overpriced. Second, buyers may infer that underpriced homes are lemons. Third, a house's relative list price may be a signal of seller type, such as an unwillingness to negotiate (Albrecht et al., 2013). Fourth, homes with high list prices may be less likely to sell quickly and may consequently be more exposed to the tail risk of becoming a "stale" listing that sits on the market without selling (Taylor, 1999). Fifth, buyers may infer that underpriced homes have a higher effective price than their list price because their price is likely to be increased in a bidding war (Han and Strange, 2012b).

Nonetheless, concrete evidence is needed for the existence of concave demand in housing markets before it is adopted as an explanation for momentum. Consequently, this section assesses whether demand is concave by analyzing micro data on listings matched to sales outcomes for the San

[12] Gopinath and Itshoki (2010) review both the price microdata and exchange rate pass-through literatures and argue there is a collage of evidence supporting a role for strategic complementarity in wholesale prices, but not resale prices. The most direct evidence to date comes from Nakamura and Zerom (2010), who directly estimate the "super elasticity" (rate of change of the elasticity) of demand for coffee using a random coefficients structural model and find evidence for concave demand.

Francisco Bay, Los Angeles, and San Diego metropolitan areas from April 2008 to February 2013.[13]

The relevant demand curve for list-price-setting sellers is the effect of unilaterally changing a house's relative quality-adjusted list price relative on its probability of sale. Detecting a nonlinear effect is challenging because quality is poorly measured, list prices are endogenous, and market conditions vary. The principal econometric challenge is that quality differences unobserved to the econometrician lead to an estimated demand curve that is far more inelastic than the true demand curve. The analysis is also complicated by the high number of foreclosures and short sales during the period that I analyze. Short sales, which occur when a home is sold for less than the outstanding mortgage balance, are especially worrisome because they often involve lengthy negotiations between the seller and their mortgage servicer which artificially decrease the probability of sale.

To surmount these challenges, I use a non-linear instrumental variable approach that traces out the demand curve using plausibly exogenous supply-side variation in seller pricing behavior. Before explaining the econometric strategy and presenting my main estimates, I first discuss the data.

1.3.1 Data

I combine data on listings with data on housing characteristics and transactions. The details of data construction can be found in Appendix A.1. The listings data come from Altos Research, which every Friday records a snapshot of homes listed for sale on multiple listing services (MLS) from several publicly available web sites and records the address, MLS identifier, and list price. The housing characteristics and transactions data come from DataQuick, which collects and digitizes public records from county register of deeds and assessor offices. This data provides a rich one-time snapshot of housing characteristics from 2013 along with a detailed transaction history of each property from 1988 to 2013 that includes transaction prices, loans, buyer and seller names and characteristics, and seller distress. I limit my analysis to non-partial transactions of single-family existing homes as categorized by DataQuick.

I match the listings data to a unique DataQuick property ID. To account for homes being de-listed and re-listed, listings are counted as contiguous if the same house is re-listed within 90 days and there is not an intervening foreclosure. If a matched home sells within 12 months of the final

[13] These metro areas were selected because both the listings and transactions data providers are based in California, so the matched dataset for these areas is of high quality and spans a longer time period.

listing date, it is counted as a sale, and otherwise it is a withdrawal. The matched data includes 83 percent of single-family transactions in the Los Angeles area and 73 percent in the San Diego and San Francisco Bay areas. It does not account for all transactions due to three factors: a small fraction of homes (under 10%) are not listed on the MLS, some homes that are listed in the MLS contain typos or incomplete addresses that preclude matching to the transactions data, and Altos Research's coverage is incomplete in a few peripheral parts of each metropolitan area.

I limit the data to homes listed between April 2008 and February 2013.[14] I drop cases in which a home has been rebuilt or significantly improved since the transaction, the transaction price is below $10,000, or a previous sale occurred within 90 days. I exclude ZIP codes with fewer than 500 repeat sales between 1988 and 2013 because my empirical approach requires that I calculate a local house price index. These restrictions eliminate approximately five percent of listings.

The final data set consists of 665,560 listings leading to 467,456 transactions. I focus on the 431,830 listings leading to 318,842 transactions with an observed prior transaction, and my IV procedure is limited to a more restricted sample described below. Table 2 provides summary statistics for several different subsamples.

1.3.2 Empirical Approach

Econometric Model Before presenting the empirical approach, I introduce an econometric framework for how changes in list price around a quality-adjusted average price affect probability of sale. Each possible sequence of list prices is associated with a distribution of time to sale. To simplify the analysis, the unit of observation is a listing associated with an initial log list price, p. I work with a summary statistic of the time to sale distribution, d, which in the main text is an indicator for whether the house sells within 13 weeks, with a withdrawal counting as a non-sale. I vary the horizon and use time to sale for the subset of listings that sell in robustness checks. The data consist of homes, denoted with a subscript h, from markets defined by a location ℓ (a ZIP code in the data) and time period t (a quarter in the data).

[14] The Altos data begins in October 2007 and ends in May 2013. I allow a six month burn-in so I can properly identify new listings, although the results are not substantially changed by including October 2007 to March 2008 listings. I drop listings that are still active on May 17, 2013, the last day for which I have data. I also drop listings that begin less than 90 days before the listing data ends so I can properly identify whether a home is re-listed within 90 days and whether a home is sold within six months. The Altos data for San Diego is missing addresses until August 2008, so listings that begin prior to that date are dropped. The match rate for the San Francisco Bay area falls substantially beginning in June 2012, so I drop Bay area listings that begin subsequent to that point.

Table 2: Summary Statistics For Listings Micro Data

Sample	All All	Prior Trans All	IV All	All Transactions	Prior Trans Transactions	IV Transactions
Transaction	70.20%	73.80%	66.80%	100%	100%	100%
Prior Transaction	64.90%	100%	100%	68.20%	100%	100%
REO	20.50%	24.90%	0%	26.70%	31.90%	0%
Short Sales	20.60%	24.20%	0%	20.20%	23.70%	0%
Positive Appreciation Since Purchase		43.00%	100%		42.30%	100%
Initial List Price	$642,072	$586,010	$817,797	$581,059	$541,682	$789,897
Transaction Price				$534,886	$ 497,901	$731,757
Weeks on Market				15.07	15.69	12.39
Sold Within 13 Wks	43.30%	44.10%	46.80%	61.70%	59.70%	70.10%
Beds	3.28	3.24	3.31	3.27	3.23	3.30
Baths	2.19	2.12	2.28	2.15	2.10	2.26
Square Feet	1,810.10	1,722.10	1,910.40	1,762.40	1,694.30	1,887.50
N	665,560	431,830	111,293	467,456	318,842	74,299

Notes: Data covers listings between April 2008 and February 2013 in the San Francisco Bay, Los Angeles, and San Diego areas as described in Appendix A.1. REOs are sales of foreclosed homes and foreclosure auctions. Short sales include cases in which the transaction price is less than the amount outstanding on the loan and withdrawals that are subsequently foreclosed on in the next two years. Appreciation since purchase is based on the ZIP code repeat-sales price index described in Appendix A.1.

I am interested in the impact of quality-adjusted list price relative to the average quality-adjusted list price in the market on probability of sale.[15] The quality-adjusted average list price $\tilde{p}_{h\ell t}$ has two additive components: the average log list price in location ℓ at time t, represented by a fixed effect $\xi_{\ell t}$, and quality $q_{h\ell t}$ that is only partially observable to the econometrician:

$$\tilde{p}_{h\ell t} = \xi_{\ell t} + q_{h\ell t}. \tag{2}$$

In a Walrasian world, there would be no variation in $p_{h\ell t} - \tilde{p}_{h\ell t}$ because sellers would all price homes at $\tilde{p}_{h\ell t}$ understanding that homes priced above $\tilde{p}_{h\ell t}$ would not sell and that pricing below $\tilde{p}_{h\ell t}$ leaves money on the table. In the housing market, however, there are search frictions and substantial amounts of idiosyncratic preference that cause demand to be a downward-sloping function of $p_{h\ell t} - \tilde{p}_{h\ell t}$, which can be thought of as the seller's relative markup. Variation in the relative markup

[15] While I focus on list prices, it is important to test the robustness of the results to using transaction prices to ensure that bargaining or price wars that occur after a list price is chosen do not undo any concavity in list price. Appendix A.3 shows all results are robust to using transaction prices.

represents differences in sellers' outside options due to factors like liquidity.

Formally, I model the probability of sale $d_{h\ell t}$ as:

$$d_{h\ell t} = g\left(p_{h\ell t} - \tilde{p}_{h\ell t}\right) + \psi_{\ell t} + \varepsilon_{h\ell t}. \tag{3}$$

The demand curve in relative price $g(\cdot)$ is assumed to be invariant across markets defined by a location and time net of an additive fixed effect $\psi_{\ell t}$ that represents local market conditions. $\varepsilon_{h\ell t}$ is an error term that represents luck in finding a buyer and is assumed to be independent of the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$.[16]

If $\tilde{p}_{h\ell t}$ were observable, one could directly estimate (3) by approximating $g(\cdot)$ with a flexible function and using ordinary least squares or by using non-parametric regression. However, observable measures of quality are imperfect, so quality $q_{h\ell t}$ likely has a component that is unobserved to the econometrician. I consequently model quality as a linear function of observed measures of quality $X_{h\ell t}$ and quality unobserved by the econometrician $u_{h\ell t}$:

$$q_{h\ell t} = \beta X_{h\ell t} + u_{h\ell t}. \tag{4}$$

I include two measures of each house's value at listing as quality measures in $X_{h\ell t}$: a repeat-sales predicted price equal to the price the last time the house sold converted to today's prices using a repeat-sales house price index and a predicted price from a hedonic index that values the house based on its characteristics.[17] To construct the repeat-sales predicted price, I first estimate interval-weighted geometric repeat-sales house price index for each ZIP code as in Case and Shiller (1989). The log index for a given time period is a time dummy in a regression of log house price on house and time fixed effects. The log predicted price $\hat{p}_{h\ell t}^{repeat}$ at time t for a house h in location ℓ that sold for $P_{h\ell\tau}$ at time τ is equal to $\log\left(P_{h\ell\tau}\frac{\phi_{\ell t}}{\phi_{\ell\tau}}\right)$, where $\phi_{\ell t}$ is the ZIP code repeat-sales index at time t. To construct the hedonic predicted price, I estimate a hedonic house price index for each

[16] Demand shocks like $\varepsilon_{h\ell t}$ traditionally cause an endogeneity problem because they are correlated with price. However, here the variable of interest is relative price, so the effect of demand shocks on average price levels is absorbed into $\xi_{\ell t}$. Similarly, the effect of prices on aggregate demand is absorbed into $\psi_{\ell t}$. It is thus natural to assume that $\varepsilon_{h\ell t}$ is independent of the relative markup in this framework.

[17] The inclusion of a predicted price to estimate the effect of a "markup" on probability of sale builds on Yavas and Yang (1995). More broadly, my empirical question and approach are similar to a real estate literature that seeks to assess the impact of list price on time on the market (Kang and Gardner, 1989; Knight 2002; Anglin et al. 2003; Haurin et al., 2010). This literature has not focused on nonlinearity, in part because of small sample sizes.

ZIP code using a third order polynomial in age, log square feet, bedrooms, and bathrooms for the hedonic factor. The predicted log price $\hat{p}_t^{hedonic}$ is the sum of a house's hedonic value as implied by a regression and the fixed effect in the regression for a given time period. The construction of both indices follows practices common in the literature and is detailed in Appendix A.1. I include both predicted prices in $X_{h\ell t}$ because each approach has its virtues (Meese and Wallace, 1997).[18] In Appendix A.3, I show the results are robust to modeling quality as a more flexible function of the predicted prices and to including other observables in $X_{h\ell t}$.

Combining (2) and (4), the reference price $\tilde{p}_{h\ell t}$ can be written as:

$$\tilde{p}_{h\ell t} = \xi_{\ell t} + \beta X_{h\ell t} + u_{h\ell t} \tag{5}$$

where again $\xi_{\ell t}$ is a fixed effect that represents the average price in location ℓ at time t and $u_{h\ell t}$ is unobserved quality.

Instrument To identify the demand curve $g(\cdot)$ in the presence of unobserved quality, I use plausibly exogenous supply-side variation in the list price due to the liquidity needs of sellers. Sellers face a trade-off between selling at a higher price and selling faster. Sellers with less liquidity and consequently a higher marginal utility of cash on hand choose a higher list price and longer time on the market. A proxy for liquidity that is orthogonal to unobserved quality and seller patience can is thus an instrument for list price.

The proxy for liquidity that I use is the equity a seller extracts from their sale. Housing is a large component of household wealth, and many sellers use the equity they extract from sale for the down payment on their next home (Stein, 1995). This increases the marginal utility of cash on hand for sellers who extract very little equity from their house because each additional dollar of equity they extract can be leveraged to buy a substantially better house. The marginal utility of cash is lower for sellers extracting substantial equity because their purchasing power is limited more by their creditworthiness and overall budget than the cash they have on hand. Consequently, homeowners with lower equity positions set higher list prices and sell their houses at higher prices

[18] The hedonic approach uses a limited set of characteristics and assumes that their valuation over time is constant because I have only a single snapshot of characteristics, but it uses all sales. Repeat sales controls for home fixed effects but only uses a subset of the data and assumes that house quality is constant and that the set of houses trading at any given time is representative.

(Genesove and Mayer, 1997; Genesove and Mayer, 2001).

Because financing and refinancing decisions make the equity of sellers endogenous, I use as my instrument the log of appreciation in the ZIP repeat-sales house price index since purchase $z_{h\ell t} = \log\left(\frac{\phi_{\ell t}}{\phi_{\ell \tau}}\right)$, where ϕ is the repeat-sales house price index, t is the period of listing, and τ is the period of previous sale.[19] This would be isomorphic to equity if all homeowners took out an identical mortgage and did not refinance. The instrument thus compares sellers who purchase identical homes with identical mortgages but who have different amounts of cash on hand to make their next down payment because one seller's home appreciated more in value than the other's.

If variation in seller liquidity represented by $z_{h\ell t}$ is independent of unobserved quality and is the only source of variation in price conditional on quality and average price, $z_{h\ell t}$ can be used as an instrument to trace out the demand curve $g(\cdot)$. Because existing evidence shows that the effect of equity is non-linear and strongest for sellers with low equity (Genesove and Mayer, 1997), I let $z_{h\ell t}$ affect price through a flexible function $f(\cdot)$. Formally, $g(\cdot)$ is identified if:

Condition 1

$$z_{h\ell t} \perp\!\!\!\perp (u_{h\ell t}, \varepsilon_{h\ell t})$$

and

$$\begin{aligned} p_{h\ell t} &= f(z_{h\ell t}) + \tilde{p}_{h\ell t} \\ &= f(z_{h\ell t}) + \xi_{\ell t} + \beta X_{h\ell t} + u_{h\ell t}. \end{aligned} \tag{6}$$

The first half of Condition 1 is an exclusion restriction that requires that appreciation since purchase have no direct effect on the outcome, either through fortune in finding a buyer $\varepsilon_{h\ell t}$ in equation (3) or through unobserved quality $u_{h\ell t}$. If this is the case, $z_{h\ell t}$ only affects probability of sale through the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$. Because I use ZIP $\times$ quarter of listing fixed effects, the variation in $z_{h\ell t}$ comes from sellers who sell at the same time in the same market but purchased at different points in the cycle. Condition 1 can thus be interpreted as requiring that unobserved quality be independent of when the seller purchased.

This assumption is difficult to test because I only have a few years of listings data, so flexibly

[19] Here $z_{h\ell t}$ is a measure of liquidity, whereas when multiplied by the previous price $P_{h\ell\tau}$ in $X_{h\ell t}$ it is used to convert the previous price to present values and constrained to have the same coefficient as the previous price $P_{h\ell t}$.

controlling for when a seller bought weakens the effect of the instrument on price in equation (6) and widens the confidence intervals to the point that any curvature is not statistically significant. Nonetheless, I evaluate the identification assumption in four ways as documented in Appendix A.3. First, I vary the observable measures of quality. Second, I include including a linear time trend in date of purchase or time since purchase. Third, I limit the sample to sellers who purchased prior to 2004 and again include a linear time trend, eliminating variation from sellers who purchased near the peak of the bubble or during the bust. In all three cases, the results remain robust. Finally, I show that the shape of the estimated demand curve is similar for IV and OLS, although OLS results in a more inelastic demand curve due to the bias created by unobserved quality. While these tests assuage some concerns, if homes with very low appreciation since purchase are of substantially lower unobserved quality despite their higher average list price, my identification strategy would overestimate the true amount of curvature in the data.[20]

I focus on sellers for whom the exogenous variation is cleanest and consequently exclude three groups. First, many individuals who have had negative appreciation since purchase are not the claimant on the residual equity in their homes—their mortgage lender is. For these individuals, appreciation since purchase is directly related to how far underwater they are, which in turn affects the foreclosure and short sale processes of the mortgage lender or servicer. Because I am interested in market processes, I exclude short sales, withdrawals that are subsequently foreclosed upon, and individuals who have had negative appreciation since purchase from the analysis. Second, mortgage servicers and government-sponsored enterprises selling foreclosed homes have no reason to be sensitive to the amount of appreciation since the foreclosed-upon homeowner purchased and are dropped. Finally, investors who purchase, improve, and flip homes typically have a low appreciation in their ZIP code since purchase but improve the quality of the house in unobservable ways. To minimize the effect of investors, I exclude sellers who previously purchased with all cash, a hallmark of investors.

The second part of Condition 1 requires that liquidity embodied in $z_{h\ell t}$ is the only reason for variation in $p_{h\ell t} - \tilde{p}_{h\ell t}$. This is a strong assumption because there may be components of liquidity that are unobserved or other reasons that homeowners list their house at a price different from

[20] One concern is that sellers with higher appreciation since purchase improve their house in unobservable ways with their home equity. However, this would create a positive relationship between price and appreciation since purchase while I find a strong negative relationship.

$\tilde{p}_{h\ell t}$, such as heterogeneity in discount rates. If the second part of the condition did not hold, then the estimates would be biased because the true $p_{h\ell t} - \tilde{p}_{h\ell t}$ would equal $f(z_{h\ell t}) + \zeta_{h\ell t}$, and the unobserved error $\zeta_{h\ell t}$ enters $g(\cdot)$ nonlinearly.

However, if other sources of variation in the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$ are independent of the variation induced by the instrument, the error in $p_{h\ell t} - \tilde{p}_{h\ell t}$ would not cause spurious concavity. Intuitively, noise in $p_{h\ell t} - \tilde{p}_{h\ell t}$ would cause the observed probability of sale at each observed $p_{h\ell t} - \tilde{p}_{h\ell t}$ to be an average of the probabilities of sale at true $p_{h\ell t} - \tilde{p}_{h\ell t}$s that are on average evenly scrambled. Consequently, although the slope may be biased, the curvature of a monotonically-decreasing demand curve is preserved. An analytical result can be obtained if the true $g(\cdot)$ is a cubic regression function as in Hausman et al. (1991):

Lemma 2 *Consider the econometric model described by (3) and (5) and suppose that:*

$$z_{h\ell t} \perp\!\!\!\perp (u_{h\ell t}, \varepsilon_{h\ell t}), \tag{7}$$

$$p_{h\ell t} = f(z_{h\ell t}) + \zeta_{h\ell t} + \tilde{p}_{h\ell t}, \tag{8}$$

$\zeta_{h\ell t} \perp\!\!\!\perp f(z_{h\ell t})$, and the true regression function $g(\cdot)$ is a third-order polynomial. Then estimating $g(\cdot)$ assuming that $p_{h\ell t} = f(z_{h\ell t}) + \tilde{p}_{h\ell t}$ yields the true coefficients of the second- and third-order terms in $g(\cdot)$.

Proof. See Appendix A.3. ■

While a special case, Lemma 2 makes clear that the bias in the estimated concavity is minimal if $\zeta_{h\ell t} \perp\!\!\!\perp f(z_{h\ell t})$. Appendix A.3.5 shows more generally using Monte Carlo simulation that if $\zeta_{h\ell t} \perp\!\!\!\perp f(z_{h\ell t})$, the degree of concavity is if anything under-estimated.

However, spurious concavity is possible if other sources of variation in the relative markup are correlated with the instrument. Specifically, Appendix A.3.5 presents Monte Carlo simulations that show that if the instrument captures most of the variation in the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$ at low levels of appreciation since purchase but very little of the variation at high levels of appreciation since purchase, spurious concavity is generated because the slope is attenuated for low relative markups but not high relative markups. However, quantitatively an extreme amount of unobserved variation in the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$ is necessary to spuriously generate the amount of

concavity in the data.

Estimation Under Condition 1, $p_{h\ell t} - \tilde{p}_{h\ell t} = f(z_{h\ell t})$, and $g(\cdot)$ can be estimated by a two-step procedure that first estimates equation (6) and then uses the predicted $f(z_{h\ell t})$ as $p_{h\ell t} - \tilde{p}_{h\ell t}$ to estimate equation (3). Both equations are estimated by OLS, and in the main text I weight the specifications by the inverse standard deviation of the error in the repeat-sales index to account for the reduced precision of the predicted prices in areas with fewer transactions. I use a third-order polynomial for $f(\cdot)$. Appendix A.3 shows that the results are robust to the order of the polynomial used for $f(\cdot)$.

I approximate $g(\cdot)$ in three ways. First, I use a three-part spline in the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$, with the knot points spaced so that each segment includes one-third of the data, which allows for a statistical of nonlinearity. I calculate standard errors by block bootstrapping the entire procedure and clustering on 35 units defined by the first three digits of the ZIP code (ZIP-3).[21] Second, to visualize the data, I construct a binned scatter plot, which bins the data into 25 equally-sized groups of the log list price relative to the reference price, $p_{h\ell t} - \tilde{p}_{h\ell t}$, and, for each bin, plots the mean of $p_{h\ell t} - \tilde{p}_{h\ell t}$ against the mean of the probability of sale net of the average probability of sale in the market, $d_{h\ell t} - \psi_{\ell t}$. This approximates $g(\cdot)$ using indicator variables for the 25 bins of $p_{h\ell t} - \tilde{p}_{h\ell t}$, as detailed in Appendix A.3. Third, I use a third-order polynomial to approximate $g(\cdot)$ and plot the estimated polynomial and 95 percent confidence bands with the binned scatter plot.

There may be small-sample bias introduced into the estimation if $g(\cdot)$ is non-linear and the fixed effects $\xi_{\ell t}$ are imprecisely estimated with a small number of homes in a ZIP-quarter cell.[22] Appendix A.3 shows that the results are not substantially changed by limiting the sample to fixed effect cells with at least 15 homes. Because the error in the estimated fixed effects is likely minimal for these cells, this suggests that imprecision in the estimated fixed effects is not driving the results.[23]

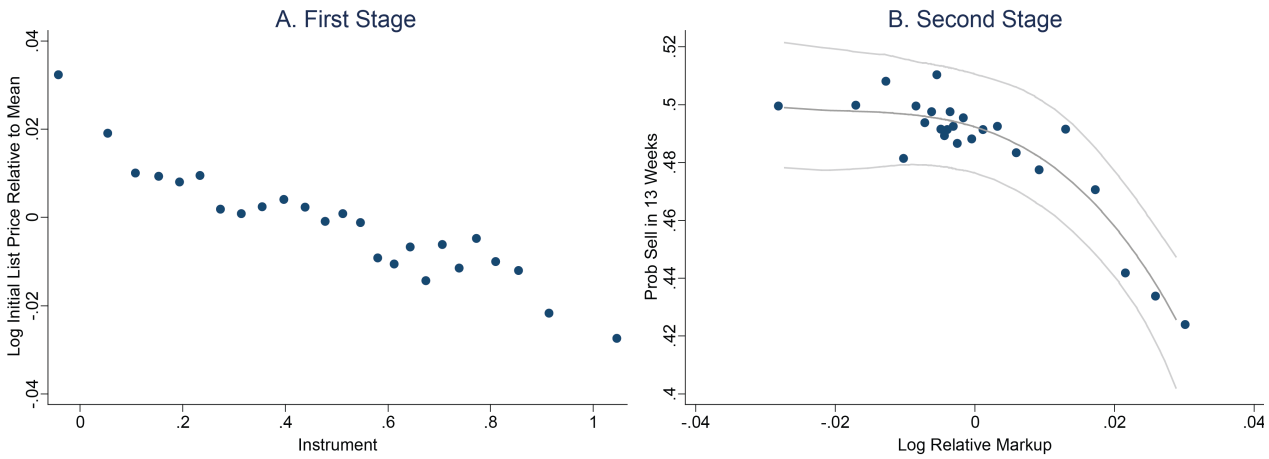

Figure 5: Instrumental Variable Estimates of the Effect of List Price on Probability of Sale

Notes: Panel B shows a binned scatter plot of the probability of sale within 13 weeks net of fixed effects (with the average probability of sale within 13 weeks added in) against the estimated log relative markup $p - \tilde{p}$. It also shows an overlaid cubic fit of the relationship, as in equation (3). To create the figure, a first stage regression of the log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and repeat sales and hedonic log predicted prices, as in (6), is estimated by OLS. The predicted value of the polynomial of the instrument is used as the relative markup. The figure splits the data into 25 equally-sized bins of this estimated relative markup and plots the mean of the estimated relative markup against the mean of the probability of sale within 13 weeks net of fixed effects for each bin, as detailed in Appendix A.3. Before binning, the 1st and 99th percentiles of the log sale price residual and any observations fully absorbed by fixed effects are dropped. The entire procedure is weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. The sample is limited to the IV subsample of homes that are not sales of foreclosures or short sales, sales of homes with negative appreciation since the seller purchased, or sales by investors who previously purchased with all cash. The grey bands indicate a pointwise 95-percent confidence interval for the cubic fit created by block bootstrapping the entire procedure on 35 ZIP-3 clusters. Panel A shows the first stage relationship between the instrument and log initial list price in equation (6) by residualizing the instrument and the log initial list price against the two predicted prices and fixed effects, binning the data into 25 equally-sized bins of the instrument residual, and plotting the mean of the instrument residual against the mean of the log initial list price residual for each bin. N = 111,293 observations prior to dropping the 1st and 99th percentiles and unique zip-quarter cells.

1.3.3 Results

Figure 5 shows the resulting first and second stage binned scatter plots. As shown in panel A, the instrument induces a small amount of variation in the list price set by sellers.[24] This is the variation

[21] I do not bootstrap the estimation of the house price indices and the predicted prices. This may add noise through a generated regressor problem (Murphy and Topel, 1985).

[22] There are 9,200 fixed effects. Less than half a percent of the data is unused because there is only a single house sold in a ZIP-quarter cell.

[23] An alternative approach is to use a random effects estimator, which I am implementing in ongoing work.

[24] Genesove and Mayer (1997) find that a house with 100 percent loan-to-value ratio is on average listed at a price four percent higher than a home with an 80 percent loan-to-value ratio. Subsequent work (Genesove and Mayer, 2001) finds slightly smaller numbers conditioning on whether a seller has experienced a nominal loss. Nonetheless, the similarity between their four percent figure and the amount of variation induced by the instrument in my first stage is reassuring.

Table 3: The Effect of List Price on Probability of Sale: Regression Results

Panel A: Ordinary Least Squares				
Dependent Var:	Sell Within 13 Weeks			
Sample:	All Listings With Prior Observation			
	(431,830 obs, 420,820 After Dropping 1st and 99th % and Cells With One Obs)			
Controls:	ZIP × Quarter × Distress FE, Repeat and Hedonic Predicted Price			
	Lowest Tercile	Middle Tercile	Highest Tercile	High - Low
Coefficient on List Price	0.161***	-0.500***	-0.483***	-0.643***
Residual Spline	(0.031)	(0.091)	(0.039)	(0.056)
Bootstrapped 95% CI				[-0.767,-0.555]
Panel B: Instrumental Variable				
Dependent Var:	Sell Within 13 Weeks			
Sample:	Listings With Prior Obs, excluding REO, Short Sales, Investors, Neg Appreciation			
	(111,293 obs,108,696 After Dropping 1st and 99th % and Cells With One Obs)			
Controls:	ZIP × Quarter FE, Repeat and Hedonic Predicted Price			
Instrument:	Appreciation Since Purchase			
	Lowest Tercile	Middle Tercile	Highest Tercile	High - Low
Coefficient on List Price	-0.320	0.261	-2.327***	-2.007***
Residual Spline	(0.334)	(1.651)	(0.616)	(0.588)
Bootstrapped 95% CI				[-3.577,-1.293]

Notes: * p < 0.05, ** p<0.01, *** p<0.001. Each row shows regression coefficients when g(.) in equation (3) is approximated using a three-segment linear spline with an equal fraction of the data in each segment. This relationship represents the effect of the log relative markup on the probability of sale within 13 weeks. In the IV panel, a first stage regression of log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and log predicted price using both a repeat-sales and a hedonic methodology, as in (6) is estimated by OLS. The predicted value of the polynomial of the instrument is used as the relative markup in equation (3), which is estimated by OLS. The sample is restricted to non-REOs, non-short sales, properties with positive appreciation since purchase, and properties not previously purchased with all cash (investors). In the OLS panel, quality is assumed to be perfectly measured by the hedonic and repeat-sales predicted prices and have no unobserved component. OLS thus regresses log list price on fixed effects and the predicted prices and uses the residual as the estimated relative markup into equation (3), as described in Appendix A.3. OLS uses the full set of listings with a previous observed transaction, so to prevent distressed sales from biasing the results, the fixed effects are at the quarter of initial listing x ZIP x distress status level. Distress status corresponds to three groups: normal sales, REOs (sales of foreclosed homes and foreclosure auctions), and short sales (cases where the transaction price is less than the amount outstanding on the loan and withdrawals that are subsequently foreclosed on in the next two years). Both procedures are weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. Before creating the spline, the 99th and 1st percentiles of the relative markup are dropped, as are any observations fully absorbed by fixed effects. In addition to the regression coefficients, the difference between the highest and lowest tercile of the spline is reported. Standard errors and the 95 percent confidence interval for the difference between the first and third terciles are computed by block bootstrapping the entire procedure on 35 ZIP-3 clusters.

I use to identify the shape of demand. The first stage is strong with a joint F statistic for the third order polynomial of the instrument in (6) of 128. Panel B shows that a clear concave relationship

is visible in the second stage, with very inelastic demand for relatively low priced homes and elastic demand for relatively high priced homes. This curvature is also visible in the cubic polynomial fit.[25] Table 3 shows regression results when $g(\cdot)$ is approximated by a three-part spline. Panel B shows the IV results. The concavity visible in Figure 5 is apparent, with the highest tercile having a slope that is seven times the lowest tercile. The difference between the highest and lowest tercile slopes is statistically significant.

As a point of comparison, Panel A shows OLS results for the full sample of homes with a prior observed transaction. The fixed effects are at the ZIP $\times$ quarter $\times$ REO seller $\times$ short seller level to prevent distressed sales from biasing the results. OLS assumes away unobserved quality and should be positively biased if $\tilde{p}_{h\ell t}$ is positively correlated with $p_{h\ell t}$ due to omitted unobserved quality. This is the case: the estimated demand curve is more elastic for IV than OLS. In fact, the OLS bias is strong enough that the demand curve slopes significantly upward in the lowest tercile. Nonetheless, a clear pattern of concavity is apparent in the OLS results. Appendix A.3 shows that OLS looks similar on the limited IV sample.

The highest tercile IV estimates imply that raising one's price by one percent reduces the probability of sale within 13 weeks by approximately 2.3 percentage points on a base of 46.8 percentage points, a reduction of 5 percent. This corresponds to a one percent price hike increasing the time to sale by six to eleven days. This figure is of comparable magnitude to Carrillo (2012), who estimates a structural search model of the steady state of the housing market with multiple dimensions of heterogeneity using data from Charlottesville, Virginia from 2000 to 2002. Although we use very different empirical approaches, in a counterfactual simulation, he finds that a one percent list price increase increases time on the market by a week, while a five percent list price increase increases time on the market by a year. Carrillo also finds small reductions in time on the market from underpricing, consistent with the nonlinear relationship found here.

Appendix A.3 shows that the results are robust across geographies, time periods, and specifications, although in some cases restricting to a smaller sample leads to insignificant results. It also shows that concavity is clearly visible in the reduced-form relationship between the instrument and probability of sale. Finally, the Appendix shows the results are robust to other measures of quality

[25] Most of the curvature comes from the top quarter of the sample because the instrument has the largest effect on the small number of sellers with low appreciation since purchase and a smaller effect on sellers who have experience moderate to high appreciation.

and to using transaction prices rather than using list prices. The instrumental variable results thus provide evidence of demand concave in relative price for these three MSAs from 2008 to 2013.[26]

1.4 A Model of House Price Momentum

This section introduces an equilibrium search model with concave demand. The model includes two additional ingredients new to the housing search literature. First, because concave demand only amplifies existing price insensitivity, I introduce variants of the model with two separate sources of insensitivity: staggered pricing as in Taylor (1980) and a small number of backward-looking rule-of-thumb sellers as in Haltiwanger and Waldman (1989) and Gali and Gertler (1999).

Second, I include an endogenous entry decision for buyers and sellers so that the same model can be used to assess how the re-timing of purchases and sales in light of momentum affects housing dynamics. Entry is a form of intertemporal arbitrage that reduces the amount of momentum in the model, and with a completely elastic entry margin momentum would be eliminated (Barsky et al., 2007). Consequently, the model features some households who have to move immediately so that the entry margin is important but not strong enough to eliminate momentum.

The model builds on search models of the housing market, such as Wheaton (1990), Krainer (2001), Novy-Marx (2009), Piazzesi and Schneider (2009), Caplin and Leahy (2011), Genesove and Han (2012), Head et al. (2014), Ngai and Tenreyro (2013), Burnside et al. (2013), and Diaz and Jerez (2013). I also incorporate ideas from models with price posting with undirected search (*e.g.*, Kudoh, 2013).

I first introduce a framework that models a metropolitan area with a fixed population and housing stock. I then describe the housing market component and show how sellers set list prices. I then introduce staggered pricing and rule-of-thumb consumers. The notation used in the model is summarized in Tables 4 and 5.

1.4.1 Setting

Time is discrete and all agents are risk neutral. Agents have a discount factor of β and time t is denoted with a subscript. There is a fixed housing stock of mass one, no construction, and a fixed

[26] Aside from the tail end of my sample, this period was a depressed market. The similarity between my results and Carrillo's provide some reassurance that the results I find are not specific to the time period, but I cannot rule out that the nonlinearity would look different in a booming market.

Table 4: Notation in the Model

Variable	Description	Note
	Masses	
Pop	Total Population (Housing Stock Mass One)	
B	Endogenous Mass of Buyers	Value Fn V^b
S	Endogenous Mass of Sellers	Value Fn V^s
R	Endogenous Mass of Renters	Value Fn V^r
H	Endogenous Mass of Homeowners	Value Fn V^h
	Flow Utilities	
b	Flow Utility of Buyer (Includes search cost)	
s	Flow Utility of Seller (includes search cost)	
u	Flow Utility of Renter	Shocked Variable
h	Flow Utility of Homeowner	
	Moving Shock Probabilities	
λ^h	Prob Homeowner Gets Shock	
λ^r	Prob Renter Gets Shock	
	Costs	
c	Stochastic Cost for Homeowner to Stay in Home	$\sim C(c), U(\underline{c}, \bar{c})$
k	Stochastic Cost for Renter to Stay Renter (Negative)	$\sim K(k), U\left(\underline{k}, \bar{k}\right)$
c^*	Threshold c Above Which Homeowners Enter	Endogenous
k^*	Threshold k Above Which Renters Enter	Endogenous
	Other Parameters	
β	Discount Factor	
L	Probability Seller Leaves Metro Area	
V^0	Value Realized Upon Exiting Metro Area	

population of size Pop.[27] Each period occurs in three stages: first search and transactions occur, then flow utilities are realized, and finally mismatch shocks occur.

There are four types of homogenous agents: a mass B_t of buyers, S_t of sellers, H_t of homeowners, and R_t of renters. These agents have flow utilities (inclusive of search costs) b, s, h, and r, and value functions V_t^b, V_t^s, V_t^h, and V_t^r, respectively. Buyers and sellers are active in the housing market, which is described in the next section. The rental market, which serves as a reservoir of potential buyers, is unmodeled aside from the flow utility net of rents. I assume that each agent can own only one home, which precludes short sales and investor-owners, although I allow for the re-timing of buyer and seller entry decisions described below.

Each period with probability λ^h and λ^r, respectively, homeowners and renters receive shocks

[27] Construction is omitted for parsimony. The model best applies to areas with inelastic housing supply in which momentum is stronger, although it is also relevant to the short run in elastically supplied metro areas, in which momentum is weaker but still important. See Head et al. (2014) for a model with a construction margin.

that cause them to separate from their current house or apartment, as in Wheaton (1990). However, rather than automatically entering the housing market, the shocks cause homeowners and renters to draw a one-time cost, $c \sim C(\cdot)$ for homeowners and $k \sim K(\cdot)$ (likely negative) for renters, that can be paid to stay in their current house or apartment and receive the same flow utility as before instead of moving. Because the seller entry elasticity appears to be constant over the cycle as shown in Appendix A.5, the cost distributions are parameterized as uniform: $c \sim U(\underline{c}, \bar{c})$ and $k \sim U\left(\underline{k}, \bar{k}\right)$. This setup captures that potential movers have heterogeneous reasons to buy or sell and consequently differ in the ease with which they can re-time their transaction.

A renter who decides not to pay the cost k enters the market as a homogenous buyer. A homeowner who decides not to pay the cost c learns after making their entry decision whether they leave the MSA with probability L, in which case they become a seller and receive termination payoff V^0 for leaving, or whether they remain in the city with probability $1-L$. If they remain in the city, they simultaneously become a buyer and a homogenous seller. These two roles are assumed to be quasi-independent so that the value functions do not interact and no structure is put on whether agents buy or sell first, as in Ngai and Tenreyro (2013) and Guren and McQuade (2013).

A homeowner who draws a cost c enters the market if:

$$c \geq V_t^h - V_t^s - LV^0 - (1-L)V_t^h \equiv c_t^*. \tag{9}$$

Similarly a renter enters if

$$k \geq V_t^r - V_t^b \equiv k_t^*. \tag{10}$$

The cutoffs c_t^* and k_t^* determine the marginal buyer and seller and control their flow into the market.[28]

Because the population is constant, every time a seller leaves the city they are replaced by a new entrant. Entrants draw a cost of being a renter and decide whether to rent or buy in the same manner as a renter who just experienced a shock. The full closed system is illustrated diagrammatically in Figure 6. The laws of motion and value functions of a homeowner and renter are deferred to Appendix A.4.

[28] This setup makes two implicit assumptions for tractability. First, although individuals are heterogeneous in their motivation to move, once they enter the market they are homogenous. Second, if an individual decides not to move today, they do not make another decision about moving until they get another shock.

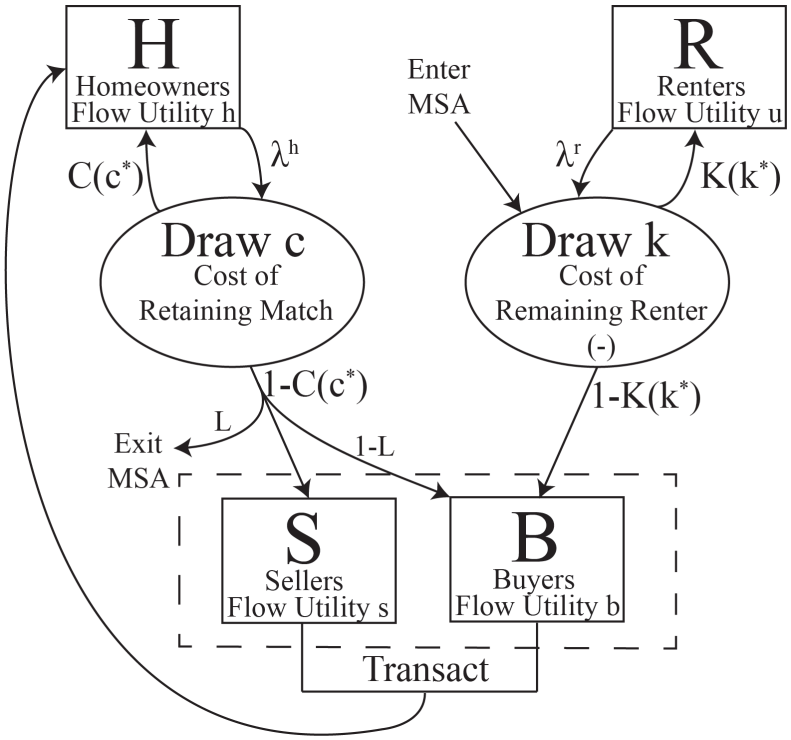

Figure 6: Schematic Representation of Closed GE System

1.4.2 The Housing Market

The search process occurs at the beginning of each period and unfolds in three stages. First, sellers post list prices $\hat{p}_t$.[29] Second, buyers search and stochastically find a single house to inspect. Third, matched buyers inspect the house. When they do so, they observe their idiosyncratic valuation for the house ε_m, which is match-specific, drawn from $F(\varepsilon_m)$ at inspection, and realized as utility at purchase. They also observe the house's permanent quality υ_h, which is mean-zero, gained by a buyer at purchase, and lost by a seller at sale. The buyer then decides whether to purchase the house or to continue searching.

I assume all sales occur at list price, or equivalently that risk neutral buyers and sellers expect that the average sale price will be an affine function of the list price.[30] This assumption is made for tractability and is not essential to the propagation mechanism. It is also less strong than it may

[29] Lester et al. (2013) show that list prices are an optimal mechanism when inspection is costly. Intuitively, a list price acts as a commitment by sellers not to waste buyers' time.

[30] This assumption restricts what can occur in bargaining or a price war. Several papers have considered the role of various types of bargaining in a framework with a list price in a steady state search model, including cases in which the list price is a price ceiling (Chen and Rosenthal, 1996; Haurin et al., 2010), price wars are possible (Han and Strange, 2013), and list price can signal seller type (Albrecht et al., 2013).

Table 5: Notation in Housing Market

Variable	Description	Note
	Utilities	
ε_m	Match-Specific One-Time Utility Benefit	$\sim F(\varepsilon)$
υ_h	Permanent House Quality	Mean Zero
	Stochastic Draws	
$\eta_{h,t}$	Noise in Observed υ_h, IID Common in Period t	$\sim G(\eta)$
	Parameters / Values	
θ	Market Tightness $= B/S$	Endogenous
$\tilde{\theta}$	Effective Market Tightness $= B/S^{visited}$	Endogenous
$q\left(\tilde{\theta}\right)$	Prob. Seller Meets Buyer (Matching Function)	Endogenous
ξ	Constant in Matching Function	
γ	Matching Function Elasticity	
Ω	Distribution of Prices	Endogenous
ε^*	Threshold ε_m for Purchase	Endogenous
μ	Threshold for Binary Signal	
	Distribution Parameters	
χ	Exponential Dist Param for $F(\varepsilon)$	
σ	Logistic Variance Param for $G(\eta)$	

first appear: although many houses do sell above or below list price, Appendix A.1.3 shows that in the merged Altos-DataQuick micro data, the modal transaction price is the list price, and the average and median differences between the list and transaction price are less than 0.01 log points and do not vary much across years.[31]

Buyer search for homes is partially directed in that buyers search only for homes that do not appear overpriced for their quality, but whether a house is overpriced for its quality is noisily observed. This directs search away from overpriced homes but preserves much of the structure of random search in which frictions prevent buyers from seeking out the lowest price house relative to quality or the house with which they have the best match. Formally, after prices are posted, buyers receive a binary signal from their real estate agent or from advertisements. The signal reveals whether a house's quality-adjusted price relative to the average quality-adjusted price is above a threshold. However, quality υ_h is subject to mean zero noise $\eta_{h,t} \sim G(\cdot)$, where $G(\cdot)$ is assumed to be constant over time. This noise, which represents how well a house is marketed in a given

[31] An important feature of the housing market is that most price changes are decreases. Consequently, the difference between the initial list price and the sale price fluctuates substantially over the cycle as homes that do not sell cut their list price. I abstract from such duration dependence to maintain a tractable state space.

period, is common to all buyers but independent and identically distributed across periods.

Buyers can inspect at most one home per period, and although they can limit their search set, it is assumed that once they do so they search randomly among homes in their search set and cannot direct their search to a particular type of home. Buyers who only observe this signal before choosing their search optimally limit their search to homes that the signal indicates are not overpriced.[32] These are homes for which the quality-adjusted price $p_t \equiv \hat{p}_t - \upsilon_h$ satisfies,

$$p_t - \eta_{h,t} - E_\Omega [p_t] \leq \mu, \tag{11}$$

where Ω is the distribution of prices. Because the signal reveals nothing else about the home, buyers cannot do better than searching randomly within homes satisfying (11). I assume that search occurs according a constant returns to scale matching function so that the number of matches can be written as a function of the number of buyers and visited sellers $m\left(B_t, S_t^{visited}\right)$. Because m is constant returns to scale, I rewrite m as a function $q\left(\tilde{\theta}_t\right)$ of the ratio of buyers to visited sellers $\tilde{\theta}_t = \frac{B_t}{S_t^{visited}} = \frac{B_t}{S_t E_\Omega[1-G(p_t - E_\Omega[p_t]-\mu)]}$. The matching function captures frictions in the search process that prevent all reasonably-priced homes and all buyers from having an inspection each period. For instance, buyers randomly allocating themselves across houses may miss a few houses, or there may not be a mutually-agreeable time for a buyer to visit a house in a given period.

After inspecting a house, buyers purchase if their surplus from doing so $V_t^h + \varepsilon_m - p - b - \beta V_t^b$ is positive. This leads to a threshold rule to buy if $\varepsilon_m > p_t + b + \beta V_{t+1}^b - V_t^h \equiv \varepsilon_t^*$ and a probability of purchase given inspection of $1 - F(\varepsilon_t^*)$.[33]

Sellers have rational expectations but set their list price before $\eta_{h,t}$ is realized and without knowing the valuation of the particular buyer who visits their house. The demand curve they face when they set their price, $d\left(p_t, \Omega_t, \tilde{\theta}_t\right)$, is the *ex-ante* probability of sale for a house with a list price p_t given a distribution of list prices Ω_t and functional market tightness $\tilde{\theta}_t$. $d(p_t, \Omega_t, \theta_t)$ can be written as the product of the probability the house satisfies (11) and is searched, $1 - G(p_t - E_\Omega[p_t] - \mu)$, the probability a house that is searched matches with a buyer, $q\left(\tilde{\theta}_t\right)$, and the

[32] This behavior is optimal if only the signal is observed. If both the signal and price are observed, one can always find a prior distribution for quality such that following the signal is optimal.

[33] Because the signal reveals no information about the house's quality υ_h, posted price $\hat{p}_t$, or match quality ε_m, the search and inspection stages are independent.

probability of purchase given inspection, $1 - F(\varepsilon_t^*)$:

$$d\left(p_t, \Omega_t, \tilde{\theta}_t\right) = q\left(\tilde{\theta}_t\right)\left(1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right)\left(1 - F\left(\varepsilon_t^*\right)\right). \tag{12}$$

I parameterize the model by assuming distributions for $F(\cdot)$ and $G(\cdot)$. Specifically, I assume that $F(\varepsilon_m)$ is an exponential distribution with parameter χ and $G\left(\eta_{h,t}\right)$ is logistic with mean zero and variance $\sigma^2\frac{\pi^2}{3}$.[34] I also assume that the matching function is Cobb-Douglas $q(\theta) = \xi\theta^{-\gamma}$, as is standard in the search literature. While these assumptions matter for precise quantitative predictions of the model, they are not necessary for the intuitions it illustrates.

This setup leads to a concave demand curve with considerable curvature in the neighborhood of the average price. At above average prices, the demand curve is dominated by whether buyers include the house in their search set, creating an elastic demand curve. At below average prices, buyers include the house in their search set with high probability and the demand curve is dominated by purchase decisions based on a trade-off between idiosyncratic match quality ε_m and price, so demand is less elastic. To illustrate this, Figure 7 shows the shapes of the probability of inspection $q(\theta_t)\left(1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right)$, the probability of purchase conditional on inspection $1 - F\left(\varepsilon_t^*\right)$, and the overall demand curve $d\left(p_t, \Omega_t, \theta_t\right)$, equal to the product of the first two panels. (Note that the axes are swapped from the traditional Marshallian supply and demand diagram in order to be consistent with the empirical analysis in Section 1.3.)

1.4.3 Flexible Price Setting

If sellers can update their list price each period, the buyer and seller value functions are equal to the value of not transacting and remaining in the market next period plus the probability of purchase or sale times the party's surplus from the transaction relative to remaining in the market.

[34] It is useful to work with distributions for which the hazard rate $\frac{f}{1-F}$ and mean excess function $E\left[x - x^* | x > x^*\right]$ have closed-form analytic solutions. The exponential distribution is particularly convenient because it has a single parameter, changes in the tail density do not drive the results, and the hazard and mean excess functions are constant, although using a Weibull or Gamma yields similar results.

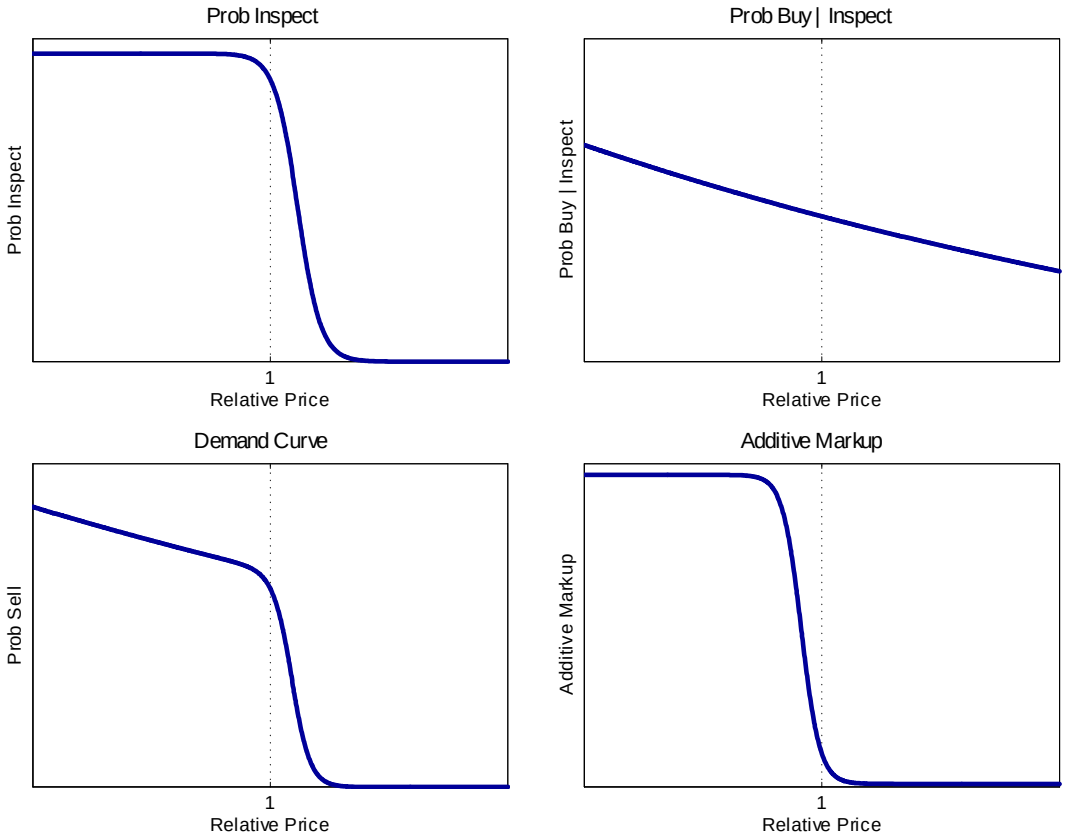

Figure 7: The Concave Demand Curve in the Model

Notes: The figures are generated using calibration described in Section 1.5. All probabilities and the additive markup are calculated assuming all other sellers are setting the steady state price and considering the effect of a unilateral deviation.

Mathematically,

$$\begin{aligned} V_t^b &= b + \beta V_{t+1}^b + \frac{d\left(p_t, \Omega_t, \tilde{\theta}\right)}{\theta_t}\left[V_t^h + \varepsilon_t^* + E\left[\varepsilon_m - \varepsilon_t^* | \varepsilon_m > \varepsilon_t^*\right] - p - b - \beta V_{t+1}^b\right] \\ &= b + \beta V_{t+1}^b + \frac{d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\chi\theta_t} \qquad (13) \\ V_t^s &= s + \beta V_{t+1}^s + \max_{p_t}\left\{d\left(p_t, \Omega_t, \tilde{\theta}_t\right)\left[p_t - s - \beta V_{t+1}^s\right]\right\}, \qquad (14) \end{aligned}$$

where (13) follows from the memoryless property of the exponential distribution for ε_m. Seller optimization implies:

Lemma 3 *The seller's optimal list price when prices can be set flexibly each period is:*

$$p = s + \beta V_{t+1}^{s} + \frac{1}{\frac{f(\varepsilon_t^*)}{1-F(\varepsilon_t^*)} + \frac{g(p_t - E_\Omega[p_t] - \mu)}{1-G(p_t - E_\Omega[p_t] - \mu)}} \tag{15}$$

$$= s + \beta V_{t+1}^{s} \frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p_t - E_\Omega[p_t] - \mu}{\sigma}\right)} + \chi}, \tag{16}$$

where the second line imposes the distributional assumptions. In a rational expectations equilibrium $p_t = E_\Omega[p_t]$*. The optimal list price is unique on an interval bounded away from* $p = \infty$*.*

Proof. See Appendix A.4.2. ■

Sellers have monopoly power due to costly search. The optimal pricing problem they solve is the same as that of a monopolist facing the demand curve d except that the marginal cost is replaced by the seller's outside option of searching again next period. The optimal pricing strategy is a markup over the outside option $s + \beta V_{t+1}^{s}$. In equation (15) it is written as an additive markup equal to the reciprocal of the semi-elasticity of demand, $\frac{-d(p_t, \Omega_t, \theta_t)}{\frac{\partial d(p_t, \Omega_t, \theta_t)}{\partial p_t}}$. The semi-elasticity, in turn, is equal to the sum of the hazard rates of the idiosyncratic preference distribution $F(\cdot)$ and the distribution of signal noise $G(\cdot)$.

This creates a strategic complementarity in price setting because the optimal price depends on relative price $p_t - E[p_t]$ through the hazard rate of the signal $G(\cdot)$. In particular, the elasticity of demand rises as relative price increases, causing the optimal markup to fall from $\frac{1}{\chi}$ to $\frac{1}{\frac{1}{\sigma}+\chi}$, as illustrated in Figure 7. The markup thus pushes sellers to set prices close to those of others. However, in a rational expectations equilibrium without additional sources of price insensitivity, all sellers choose the same list price and $p_t = E[p_t]$, so there is no relative price to affect the markup. A shock to home values thus causes list price to jump proportionally to the seller's outside option. Consequently, I introduce variants of the model with two different sources of insensitivity of prices to generate some initial momentum.

1.4.4 Source of Insensitivity 1: Staggered Price Setting

The first source of price insensitivity I consider is staggered price setting as in Taylor (1980).[35] Prices in housing markets are not constantly updated because it takes time to market a house and collect offers, and lowering the price frequently can signal that a house is of poor quality (De Wit and Van Der Klaauw, 2013).[36] While likely not the most important pricing friction in housing markets, staggered pricing has the virtue of being familiar, tractable, and quantifiable in micro data.

With N groups of sellers, denote the quality-adjusted prices p, value functions V^s, masses S, and purchase thresholds ε of a specific vintage of sellers using superscripts for the time since they last reset their price $\tau = \{0, ..., N-1\}$. The buyer's surplus from purchasing from various sellers is constant due to the memoryless property of the exponential, but the value function must be adjusted to integrate over the sellers in the market:

$$V_t^b = b + \beta E_t V_{t+1}^b + \frac{1}{\chi\theta_t} \sum_{\tau=0}^{N-1} \left[\frac{S_t^\tau}{S_t} d\left(p_t^\tau, \Omega_t, \tilde{\theta}_t\right) \right]. \tag{17}$$

The value function of a seller is similar to the frictionless case except sellers only optimize occasionally so τ superscripts are necessary:

$$V_t^{s,\tau} = s + \beta E_t V_{t+1}^{s,\tau+1} + d\left(p_t^\tau, \Omega_t, \tilde{\theta}_t\right)\left(p_t^\tau - s - \beta E_t V_{t+1}^{s,\tau+1}\right), \tag{18}$$

where $V_t^N = V_t^0$ and $d\left(p_t^\tau, \Omega_t, \tilde{\theta}_t\right)$ is as in equation (12) except ε_t^* is replaced by a separate threshold match quality $\varepsilon_t^{*,\tau}$ for each vintage of sellers.

Seller optimization implies the optimal list price is reminiscent of a Taylor (1980) or Calvo (1983) model except there is only one good to sell, so demand is replaced by that the probability the house sells in a given period:

[35] I adopt Taylor (1980) staggered pricing rather than Calvo (1983) pricing because the model includes an integral that cannot be updated iteratively in the denominator of $\tilde{\theta}_t$. Staggered pricing allows for a closed form for the integral because the price distribution has finite support.

[36] Golosov and Lucas (2008), among others, argue that models with fixed adjustment dates generate more persistence than menu cost models with state-dependent adjustment rules. As described in Section 1.5, I assume prices are fixed for two months based on data from 2008-2013, a depressed market in which sellers would have the strongest incentives to adjust their price quickly. My calibrated model thus serves as a lower bound of the frequency of price resetting one would observe in a calibrated state-dependent model.

Lemma 4 *If posted prices last N periods, the seller's optimal reset price p_t^0, where the superscript is for periods since price is set, is:*

$$p_t^0 = \frac{\sum_{\tau=0}^{N-1} \beta^\tau D_t^\tau \left(p_t^0\right) \Psi_t^\tau \varphi_t^\tau}{\sum_{\tau=0}^{N-1} \beta^\tau D_t^\tau \left(p_t^0\right) \Psi_t^\tau}, \tag{19}$$

where $D_t^j(p) = E_t\left[\prod_{\tau=0}^{j-1}\left(1 - d^\tau\left(p, \Omega_{t+\tau}, \tilde{\theta}_{t+\tau}\right)\right)\right] d\left(p, \Omega_{t+j}, \tilde{\theta}_{t+j}\right)$ *is the expected probability the house is sold τ periods after the price is set,* $\Psi_t^\tau = E_t\left[\frac{-\frac{\partial d\left(p_t, \Omega_{t+\tau}, \tilde{\theta}_{t+\tau}\right)}{\partial p_t}}{d\left(p_t, \Omega_{t+\tau}, \tilde{\theta}_{t+\tau}\right)}\right]$ *is the semi-elasticity of demand with respect to price,* $\varphi_t^\tau = s + E_t V_{t+\tau+1}^{s,\tau+1} + \frac{1}{\Psi_t^\tau}$ *is the expected optimal flexible reset price τ periods after the price is set given the expected price distribution in that period, and* $V_{t+N}^{s,N} = V_{t+N}^{s,0}$. *The optimal list price is unique on an interval bounded away from $p = \infty$ given a condition in Appendix A.4.2, which holds for all simulations considered.*

Proof. See Appendix A.4.2. ■

As is standard in staggered price models, the optimal price is a weighted average of the optimal flexible prices that are expected to prevail on the equilibrium path until the seller can reset his or her price. The weight put on the optimal flexible price in period $t + \tau$ is equal to the discounted probability of sale in period i times the semi-elasticity of demand in period i. Intuitively, the seller cares more about periods in which probability of sale is higher but also about periods in which demand is more elastic because perturbing price has a larger effect on profit.

In equilibrium, all agents behave optimally given the search technology, the noisy signal of relative price, and buyers' draw of their idiosyncratic taste when they visit a home. Laws of motion apply due to the law of large numbers. I restrict attention to symmetric equilibria. A staggered pricing equilibrium is consequently defined by:

Definition 5 *Equilibrium with N staggered groups of list-price-setting sellers is a set of prices p_t^τ, demands $d\left(p_t^\tau, \Omega, \theta\right)$, purchase cutoffs $\varepsilon_t^{*,\tau}$, and seller value functions $V_t^{s,\tau}$ for each group of sellers $\tau = \{0, ..., N-1\}$, buyer, homeowner, and renter value functions V_t^b, V_t^h, and V_t^r, entry cutoffs c_t^* and k_t^*, and stocks of each type of agent B_t, S_t^τ $\tau = \{0, ..., N-1\}$, H_t, and R_t satisfying:*

1. *Optimal reset pricing (19) and fixed pricing for non-resetters $p_t^\tau = p_{t-1}^{\tau-1}$ $\forall$ $\tau > 0$*

2. *Optimal purchasing decisions by buyers:* $\varepsilon_t^{*,\tau} = p_t^\tau + b + \beta V_{t+1}^b - V_t^h$

3. *The demand curve for each type of seller arising from optimal buyer search given the binary signal (12)*

4. *Optimal entry decisions by homeowners and renters who receive shocks (9) and (10)*

5. *The value functions for buyers (17) and each vintage of sellers (18) as well as for renters and homeowners defined in Appendix A.4*

6. *The laws of motion for all agents defined in Appendix A.4.*

Appendix A.4 shows that the model has a unique steady state that is equivalent to the frictionless case without staggered pricing ($N = 1$). A frictionless equilibrium is formally defined in the Appendix A.4.

I add a stochastic shock process to both this model and the analogous rule-of-thumb variant defined subsequently to examine their dynamic implications. The propagation mechanism for momentum does not qualitatively depend on any particular shock. However, the positive correlation between price and volume in the data implies that demand-side shocks dominate.[37] Although the particular type of demand shock introduced to the model is not important for the results, I use a shock to the flow utility of being a renter u that changes the relative value of homeownership for potential entrants. This takes a cue from Wheaton and Lee (2009), who show that changes in the frequency of transitions between renting and owning due to credit conditions are a precipitating shock for housing cycles. An example of such a shock would be a change in credit standards for new homeowners. I implement the shock by assuming that $u = \bar{u} + x$, where x is an AR(1) process understood by the forward-looking agents:

$$x_t = \rho x_{t-1} + \eta \text{ and } \eta \sim N\left(0, \sigma_\eta^2\right). \tag{20}$$

The model cannot be solved analytically, so I simulate it numerically using a log-cubic approximation pruning higher order terms as in Kim et al. (2008) implemented in Dynare (Adjemian et

[37] A positive supply-side shock to the flow value h of being a homeowner, for instance, would increase the value of homes but also induce homeowners to endogenously enter less, driving down sales volume.

al., 2013). Appendix A.6 shows that the impulse responses are similar in an exactly-solved model with a permanent and unexpected shock.

1.4.5 Source of Insensitivity 2: A Small Fraction of Rule-of-Thumb Sellers

The second source of price insensitivity I consider is a small fraction of rule-of-thumb sellers. Since Case and Shiller (1987), sellers with backward-looking expectations have been thought to play an important role in housing markets. Previous models assume that all agents have backward-looking beliefs (*e.g.*, Berkovec and Goodman, 1996), but some observers have found the notion that the majority of sellers are non-rational unpalatable given the financial importance of housing transactions for many households. Some fraction of sellers, however, may not find it worthwhile to scrutinize current market conditions due to information costs, and my model only requires a handful of backward-looking sellers because of the strategic complementarity. Consequently, I introduce a small number of rule-of-thumb sellers, as in Campbell and Mankiw (1989), and assess quantitatively what fraction of sellers is needed to be non-rational to explain the momentum in data, similar to Gali and Gertler (1999).

I assume that at all times a fraction $1-\alpha$ of sellers set their list price p_t^R rationally according to Lemma 3 and (15) but a fraction α of sellers uses a backward-looking rule of thumb to set their list price p_t^N.

The backward-looking sellers are near-rational sellers whose optimizing behavior produces a price-setting rule of thumb based on the recent price path. They are not fully rational in two ways. First, backward-looking sellers understand that a seller solves,

$$\max_{p_t} d\left(p_t, \Omega_t, \tilde{\theta}_t\right) p_t + \left(1 - d\left(p_t, \Omega_t, \tilde{\theta}_t\right)\right)\left(s + \beta V_{t+1}^s\right),$$

with first order condition,

$$p_t = s + \beta E_t V_{t+1}^s + E_t\left[\frac{-d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\frac{\partial d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\partial p_t}}\right]. \tag{21}$$

However, they do not fully understand the laws of motion and how prices and the value of being a seller evolve. Instead, they think the world is a function of a single state variable, the average

price $E[p_t]$, and can only make "simple" univariate forecasts that take the form of a first order approximation of (21) in average price and relative price:

$$p_t = s + \beta\left(\bar{V}^s_{t+1} + \pi_1 E[p_t]\right) + \bar{M} + \pi_2 E[p_t - E[p_t]], \tag{22}$$

where $\bar{V}^s$, $\bar{M}$, π_1, and π_2 are constants.

Second, backward-looking sellers only see the average prices $\bar{p}$ of houses that transact between two to four months ago and between five to seven months ago, corresponding to the lag with which reliable house price indices are released.[38] They assume that price follows a random walk with drift with both the innovations φ and the drift ζ drawn independently from mean zero normal distributions with variances σ^2_φ and σ^2_ζ. Through a standard signal extraction problem, they expect that today's price will be normally distributed with mean $E[p_t] = \bar{p}_{t-3} + E[\zeta]$, where $E[\zeta] = \frac{\sigma^2_\zeta}{\sigma^2_\zeta+\sigma^2_\varphi}(\bar{p}_{t-3} - \bar{p}_{t-6})$. Given this normal posterior, equation (22) implies $p_t = s + \beta\left(\bar{V}^s_{t+1} + \pi_1 E[p_t]\right) + \bar{M} = E[p_t]$,[39] so the backward-looking sellers follow an AR(1) rule:

$$p^N_t = \frac{p_{t-2} + p_{t-3} + p_{t-4}}{3} + \phi\left(\frac{p_{t-2} + p_{t-3} + p_{t-4}}{3} - \frac{p_{t-5} + p_{t-6} + p_{t-7}}{3}\right) \tag{23}$$

where $\phi = \frac{\sigma^2_\zeta}{\sigma^2_\zeta+\sigma^2_\varphi}$. Such an AR(1) rule is a common assumption in models with backward-looking expectations and is frequently motivated by limited knowledge, information costs, and extrapolative biases (*e.g.*, Hong and Stein, 1999; Fuster et al. 2010).[40]

I assume that the backward-looking price setters think that the variance of the innovation σ^2_ζ is a substantial share of the overall variance in price changes and consequently use a ϕ that is attenuated relative to what one would find if one ran a quarterly AR(1) in the model environment. This is consistent with Case et al. (2012), who survey home buyers for four metropolitan areas from 2003 to 2011 and show that the average predicted amount of price appreciation at a one-year horizon is approximately 43 percent of the actual amount of appreciation. An attenuated AR(1)

[38] I use three-month averages to correspond to how price indices like the closely watched Case-Shiller index are constructed and to smooth out saw-tooth patterns that emerge with non-averaged multi-period lags. A shorter AR(1) lag would require more backward-looking sellers to match the data.

[39] Specifically, $E[p_t] = s + \beta\bar{V}_s + \pi_1 E[p_t] + \bar{M}$ and so $p_t = E[p_t] + \pi_2 E[p_t - E[p_t]]$, which with a symmetric posterior for p_t implies $p_t = E[p_t]$.

[40] Coibion and Gorodnichenko also (2011) show rule-of-thumb price setters perform similarly to sticky information price setters in an estimated DSGE model.

coefficient is also consistent with psychological theories in which agents overweight and "anchor" on recent observable prices (see Barberis et al., 1998).

I make two additional assumptions for tractability and parsimony that are not crucial for the results. First, I assume that regardless of whether rational or backward-looking sellers sell faster, inflows adjust so that α of the active listings are houses owned by backward-looking sellers at all times. Second, I assume that entry occurs according to the threshold rules (9) and (10) using rational value functions. The value function of a rational seller, $V_t^{s,R}$, is the same as equation (14) for the frictionless case, while the buyer value function needs to be altered to integrate over the distribution of sellers:

$$V_t^b = b + \beta E_t V_{t+1}^b + \frac{1}{\chi\theta_t}\left[\alpha d\left(p_t^N, \Omega_t, \tilde{\theta}_t\right) + (1-\alpha)\, d\left(p_t^R, \Omega_t, \tilde{\theta}_t\right)\right]. \tag{24}$$

Given these assumptions, one can define an equilibrium with backward-looking sellers by:

Definition 6 *Equilibrium with a fraction α of backward-looking sellers is a set of prices p_t^i , demands $d\left(p_t^i, \Omega, \theta\right)$, and purchase cutoffs $\varepsilon_t^{*,i}$ for each type of seller $i \in \{N, R\}$, rational seller, buyer, homeowner, and renter value functions $V_t^{s,R}$, V_t^b, V_t^h, and V_t^r, entry cutoffs c_t^* and k_t^*, and stocks of each type of agent B_t, S_t, H_t, and R_t satisfying:*

1. *Optimal pricing for rational sellers (15) and the pricing rule (23) for backward-looking sellers*
2. *Optimal purchasing decisions by buyers: $\varepsilon_t^{*,i} = p_t^i + b + \beta V_{t+1}^b - V_t^h$*
3. *The demand curve for each type of seller arising from optimal buyer search given the binary signal (12)*
4. *Optimal entry decisions by homeowners and renters who receive shocks (9) and (10)*
5. *The value functions for buyers (24) and rational sellers (14) as well as for renters and homeowners defined in Appendix A.4*
6. *The laws of motion for all agents defined in Appendix A.4.*

The steady state of this model is the same as the staggered and frictionless models. Consequently, the staggered pricing and backward-looking models can be calibrated using the same procedure.

1.5 How Much Can Concave Demand Amplify Momentum?

To quantitatively assess the degree to which concave demand curves amplify house price momentum, this section calibrates the model to the empirical findings presented in Section 1.3 and a number of aggregate moments. Before doing so, I briefly analyze the frequency of price adjustment in the micro data to motivate the calibration of the staggered pricing variant of the model.

1.5.1 Frequency of Price Adjustment

Figure 8 shows the Kaplan-Meier survival curve for list prices of homes with an observed prior transaction in the San Francisco Bay, Los Angeles, and San Diego areas between April 2008 and February 2013. Each observation is a list price, with a sale counted as a censored observation and a price change counted as a failure. The curve thus shows the fraction of list prices that have survived a given number of weeks conditional on the house remaining on the market. The curve crosses the 50 percent threshold corresponding to the median time until a price is changed at eight weeks. Consequently, I calibrate the staggered variant of the model so that one period lasts one month, and there are two groups of sellers that alternate setting list prices that last two months.

1.5.2 Calibration and Estimation

In order to simulate the model, 21 parameters listed in Table 7 must be set. For the backward-looking variant of the model, the AR(1) coefficient in the rule of thumb ϕ and the fraction of sellers who follow it α also require numerical values. This section describes the calibration procedure and targets, with details deferred to Appendix A.5.

Three parameters control the shape of the demand curve and thus have a first-order impact on momentum: χ, the exponential parameter of the idiosyncratic quality distribution, controls the elasticity of demand for low-priced homes that are certain to be visited; σ, the logistic variance parameter of the signal, controls the elasticity of demand for high-priced homes; and μ, the threshold for being overpriced, controls where on the curve the average price lies. The other parameters affect momentum mainly through equilibrium feedbacks and largely have a second order effect on momentum. Consequently, I first estimate these three parameters from the instrumental variable micro estimates presented in Section 1.3 and then calibrate the rest of the model to match steady

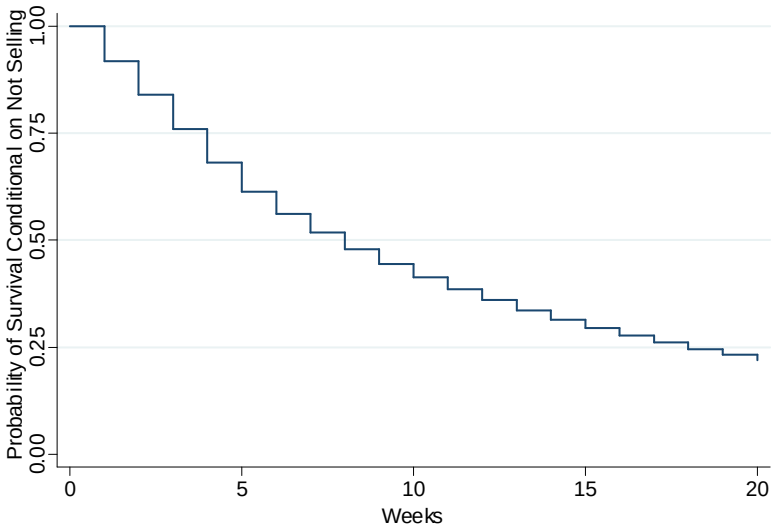

Figure 8: Kaplan-Meier Survival Curve for List Prices

Notes: The figure shows the Kaplan-Meier survival curve for list prices, where sales are treated a censored observation and a price change is treated as a failure. The curve thus corresponds to the probability of a list price surviving for a given number of weeks conditional on the property not having sold. The sample is made up of 854,547 list prices for 420,351 listings of homes with observed prior transactions in the San Francisco Bay, Los Angeles, and San Diego areas listed between April 2008 to February 2013.

state and time series aggregate moments. The calibration proceeds in two steps.

First, I estimate χ, σ, and μ to match the micro estimates. There is heterogeneity in list price in the micro estimates not in the model, and the low average probability of sale in the 2008-13 period poses a challenge because the data are not generated in a plausible steady state. To account for these features of the data, I express the probability of sale for an arbitrary distribution of prices and an arbitrary average probability of sale as functions of observable variables and the three parameters χ, σ, and μ. This allows me to approximate the model with the heterogeneity in the data out of steady state for the purposes of calibration and then conduct dynamic simulations with the heterogeneity suppressed to maintain a tractable state space.

Specifically, with my assumed functional forms, the probability of sale at the time the list price is posted can be written as:

$$\begin{aligned} d\left(p_t, \Omega_t, \tilde{\theta}_t\right) &= q\left(\tilde{\theta}_t\right)\left(1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right)\left(1 - F\left(\varepsilon_t^*\right)\right) \\ &= \kappa_t\left(1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right)\exp\left(-\chi p_t\right) \end{aligned} \tag{25}$$

The aggregate state variables factor out into a multiplicative constant, κ_t, which can be given a structural interpretation as a shift in the matching function efficiency ξ. κ_t multiplies two terms: the effect of perturbing price on the probability the house is visited $1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)$ and a term representing the buyer's trade-off between idiosyncratic quality and price $\exp\left(-\chi p_t\right)$. To simulate the probability of sale, all that is needed are χ, σ, and μ, observed prices p, the observed average price $E_\Omega\left[p\right]$, and the observed average probability of sale.

Using equation (25), I calibrate χ, σ, and μ to the IV binned scatter plot. The data is 25 ordered pairs (p_b, d_b) corresponding to the log relative markup plus the mean log price in the market and probability of sale within 13 weeks for each of 25 bins b of the distribution of the relative markup. I solve for κ_t to match the average probability of sale, and use (25) to simulate $d\left(p_b\right)$ in the model for each p_b. Because the zero point corresponding to the average price is not precisely estimated and depends on the deadline used for a listing to count as a sale, I choose the average price so that the elasticity of demand implies a monthly seller search cost of approximately \$10,000 based on evidence from Genesove and Mayer (1997) and Levitt and Syverson (2008) described in Appendix A.5.[41] In Appendix A.6.4, I evaluate the robustness of the results to this parameter by using a far smaller seller search cost. Conditional on the average price, the best fit (χ, σ, μ) is chosen to minimize the sum of squared errors $\Sigma_b w_b \left(d_b - d^{3\ month}\left(p_b\right)\right)^2$ where w_b is a Normal kernel weight to reduce the influence of outliers and $d^{3\ month}\left(\cdot\right)$ is a simulated 3-month sale probability based on (25). Figure 9 shows the IV binned scatter plot in blue Xs and the model's predicted $d\left(p_b\right)$ for the (σ, μ, χ) that minimize the distance between the model and the data in red circles. The fit suggests that the demand curve in the calibrated model captures the curvature in the data well.

The second step in the calibration is to match a number of aggregate steady state and stochastic moments given the χ, σ, and μ from the first step. I set 14 parameters to match 14 steady state moments listed in the first three panels of Table 6. These targets are either from other papers or are long-run averages for the U.S. housing market, such as the homeownership rate, the average amount of time between moves for buyers and renters, and the average time on the market for buyers and sellers. A few parameters for which data is not easily available are assumed, and the results are not sensitive to the assumed values. I also match three time series moments as indicated by the

[41] The seller search cost is likely large because of the nuisances and uncertainties involved and the need to move quickly. Another factor, highlighted by Anenberg and Bayer (2013), is the high cost of simultaneously holding two homes, which pushes households to sell quickly before buying.

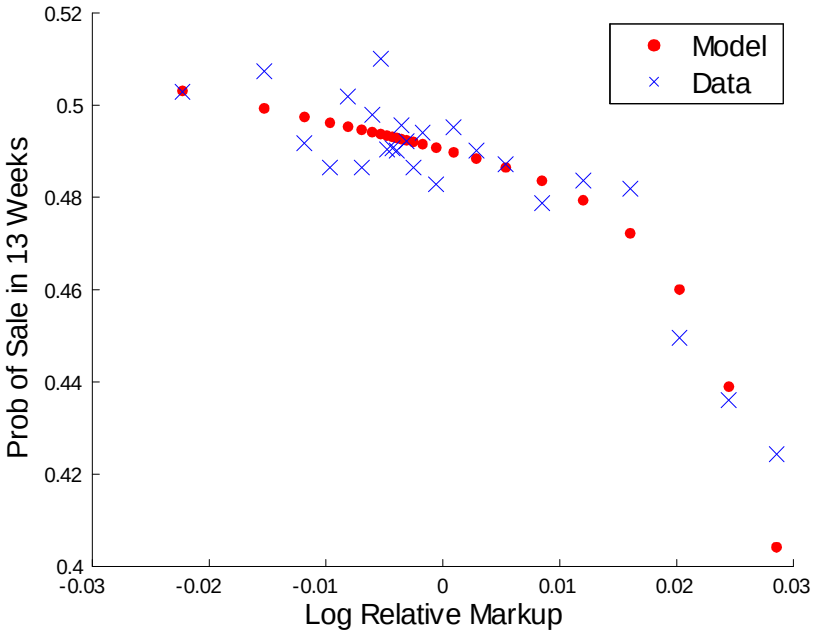

Figure 9: Model Fit Relative to Instrumental Variable Estimates

Notes: The blue Xs are the binned scatter plot from the IV specification with 2.5% of the data from each end excluded to reduce the effects of outliers. The red dots are the simulated probabilities of sale at each price level in the calibrated model.

bottom panel of Table 6. The monthly persistence of the shock is set to match the persistence of local income shocks as in Glaeser et al. (2013). The final parameters are set to match the standard deviation of annual log price changes and the elasticity of seller entry with respect to price in stochastic simulations.[42]

For the backward-looking variant of the model, I set the AR(1) coefficient ϕ to 0.4 following evidence from Case et al. (2012). Using surveys of home buyers they show that regressing realized annual house price appreciation on households' *ex-ante* beliefs yields a regression coefficient of 2.34. I use this survey evidence to calibrate the beliefs of the backward-looking sellers by dividing the approximate regression coefficient one would obtain in quarterly simulated data (approximately 0.94) by their coefficient. I adjust α and recalibrate the model until the impulse response to the renter flow utility shock matches the matches the 36 months of positively autocorrelated price changes in the AR(5) impulse response estimated on the CoreLogic national house price index in

[42] Because the stock of buyers is not observed, I cannot similarly calibrate for the buyer entry elasticity. Consequently, I assume the density of buyer entry costs is the same as the density of seller entry costs. Seller entry tends to track volume, so buyer entry cannot have a substantially different density.

Table 6: Calibration Targets

Steady State Parameter or Moment	Value	Source / Justification
	Parameters	
γ (Matching Function Elasticity)	.8	Genesove and Han (2012)
L (Prob. Stay in MSA)	.7	Anenberg and Bayer (2013)
	Aggregate Targets	
Annual Discount Rate	7%	Carrillo (2012) housing market discount rate
Time on Market for Sellers	4 Months	Approx average parameter value in literature
Time on Market for Buyers	4 Months	$\approx$ Time to sell in surveys (Genesove and Han, 2012)
Homeownership Rate	65%	Long run average, 1970s-1990s
Time in House For Owner Occupants	9 Years	American Housing Survey, 1997-2005
Time Between Moves for Renters	29 Months	American Housing Survey, 1997-2005
c^* (Cost Marginal H Pays to Avoid Move)	\$37.5k	Moving cost 5% of price (Haurin & Gill,2002)
k^* (Cost Marginal R Pays to Avoid Buying)	-\$20k	Tax benefits of owning 29 months (Poterba & Sinai, 2008)
	Assumed Values	
Time Between Shocks for Homeowners	29 Months	Same as renter
Steady State Price	\$750k	Average log price in IV sample adjusted for down market
h (Flow Utility of Homeowner)	\$7.5k	2/3 of house value from expected flow util
Prob Purchase | Inspect	0.5	So $q(\theta) \in [0, 1]$
	Time Series Moments	
SD of Annual Log Price Changes	.065	CoreLogic national HPI adjusted for CPI, 1976-2013
ρ (Monthly Persistence of AR1 Shock)	.990	Persistence of income shocks (Glaeser et al., 2013)
Price Elasticity of Seller Entry	.878	CoreLogic, Census, and NAR, 1976-2013

Section 1.2.[43]

The staggered and backward-looking variants differ minutely in their calibrated values so that each matches the volatility of price and entry elasticity in stochastic simulations, as discussed in Appendix A.5. Table 7 summarizes the calibrated parameter values for the backward-looking variant of the model.[44]

[43] An alternative approach would be to simulate data, collapse it to the quarterly level, and then estimate the same AR(5) as for the CoreLogic data. Doing so requires a fraction of rule-of-thumb price setters that is approximately ten percent higher than matching the impulse response to the model shock. A comparably higher fraction is also required to match the AR(5) without concavity. As shown in Appendix A.2, using a median price index generates an impulse response in the data that reaches its peak in two years rather than three years. My approach of calibrating the peak of the impulse response to the renter flow utility shock to the peak of the AR(5) IRF for a repeat-sales index is comparable to simulating data, estimating the AR(5) IRF, and calibrating to match the average of the repeat-sales and median price IRF peak quarters.

[44] One may argue that the flow cost of being a buyer is too large. This could be reduced without meaningfully changing the main results by relaxing several assumptions made to keep the model tractable. The buyer search cost is calibrated to a high level because the flat slope for homes priced below average implies that the exponential distribution for idiosyncratic quality has a long tail. This implies a high value of subsequent search for buyers, which is offset with a high search cost to maintain a reasonable value of being a buyer. Both using an idiosyncratic match quality distribution that is bounded above and allowing the signal to reveal more information so that buyers search

Table 7: Calibrated Parameter Values for Rule of Thumb Model

Parameter	Interpretation	Value	Parameter	Interpretation	Value
β	Monthly Discount Factor	0.994	V^0	Value of Leaving MSA	\$2,631k
γ	Matching Fn Elasticity	0.800	h	Flow Util of H	\$7.5k
ξ	Matching Fn Efficiency	0.506	u	Flow Util of R	\$3.6k
λ^h	Monthly Prob H Moving Shock	0.035	b	Flow Util of B (search cost)	-\$92.2k
λ^r	Monthly Prob R Moving Shock	0.035	s	Flow Util of S (search cost)	-\$9.8k
$\bar{c}$	Upper Bound, H Entry Cost Dist	\$463k	χ	Exponential Param for	.0023
$\underline{c}$	Lower Bound, H Entry Cost Dist	-\$1,121k		Idiosyncratic Quality Dist	
$\bar{k}$	Upper Bound, R Entry Cost Dist	\$412k	σ	Variance Param of Signal Noise	3.80
$\underline{k}$	Lower Bound, R Entry Cost Dist	-\$1,172k	μ	Threshold for Signal	10.47
Pop	Population	1.484	σ_η	SD of Innovations to AR(1) shock	0.360
L	Prob Leave MSA	0.700	ρ	Persistence of AR(1) shock	0.990
ϕ	AR(1) Param in Rule of Thumb	0.40			

Notes: The calibration is monthly. The parameters under the line are only used in the backward-looking variant of the model. The parameters for the staggered variant of the model are only minutely different and can be found in Appendix A.5.

1.5.3 Amplification of Momentum in the Calibrated Model

To assess the degree of amplification of momentum in the calibrated model, I compare the frictionless, staggered price, and backward-looking variants of the model to one another and to versions without concave demand. To do so, I examine the impulse response to the model shock to the flow utility of renters. The impulse response is computed as the average difference between two sets of simulations that use the same sequence of random shocks except for one period in which an additional standard deviation shock is added. I contrast the model impulse responses with the impulse response to a one standard deviation price shock to the quarterly CoreLogic national house price index estimated from an AR(5), as in Section 1.2.

Figure 10 shows the resulting simulations alongside the AR(5) impulse response. The figure shows that the strategic complementarity created by concave demand substantially amplifies both sources of price insensitivity.

Panel A compares a frictionless model with concave demand to staggered price models with and without concave demand, in dotted red, solid blue, and dashed green, respectively.[45] Without both

houses they expect to like would reduce the calibrated buyer search cost substantially. Ongoing work to adjust for bias in the slope of the micro estimates due to measurement error in the true relative markup, as discussed in Section 1.3, may also result in a smaller calibrated buyer search cost.

[45] The non-concave impulse response depends on the semi-elasticity of the non-concave demand function. For the main text, I assume a semi-elasticity that equal to the steady-state semi-elasticity at the average price in the concave

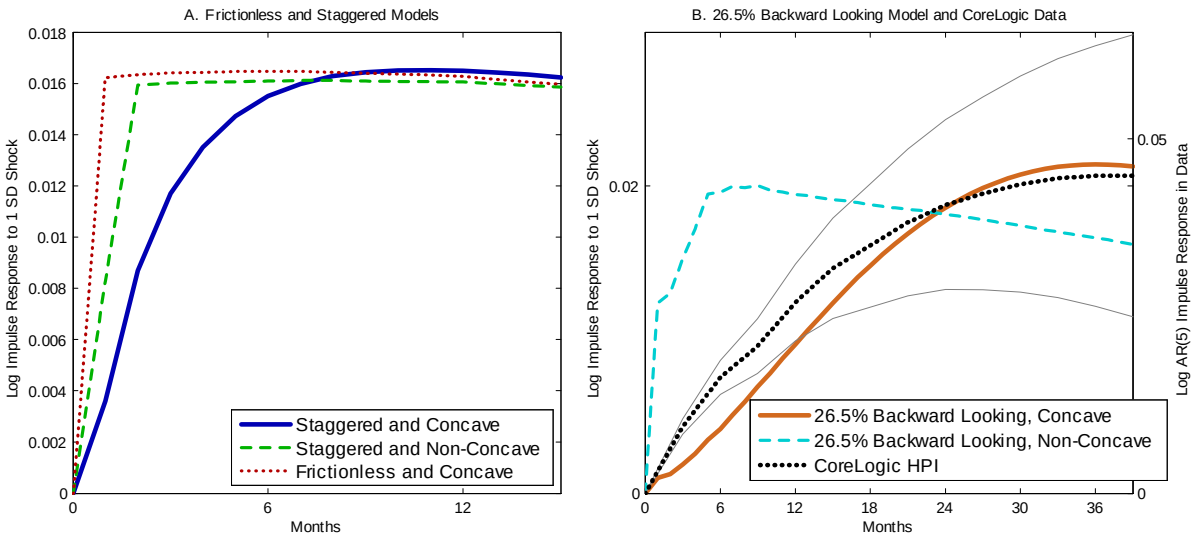

Figure 10: Price Impulse Response Functions: Model and Data

Notes: Panel A shows the impulse responses to a one standard deviation negative shock to the flow utility of renting in the frictionless model with concave demand, the staggered model with concave demand, and the staggered model without concave demand. For the model without concave demand, the threshold for being overpriced μ is raised to a level that is never reached, the slope of the demand curve is adjusted to the steady-state slope at the average price in the concave model, the model is recalibrated, and the standard deviation of the stochastic shock is adjusted so that the impulse response is even with the frictionless and concave impulse response after a year. Panel B shows the impulse responses to a one standard deviation shock to the flow utility of renting in the backward-looking model with and without concavity. For the model without concave demand, the threshold for being overpriced μ is raised to a level that is never reached, the slope of the demand curve is adjusted to the steady-state slope at the average price in the concave model, and the model is recalibrated. Also shown in panel B in the dotted black line and with grey 95% confidence intervals and on the right axis is the impulse response to a one standard deviation price shock estimated from a quarterly AR(5) for the seasonally and CPI adjusted CoreLogic national house price index for 1976-2013, as in Figure 1. Simulated impulse responses are calculated by differencing two simulations of the model from periods 100 to 150, both of which use identical random shocks except in period 101 in which a one standard deviation negative draw is added to the random sequence, and then computing the average difference over 100 simulations.

concave demand and staggering, reset prices jump on impact and reach a convergent path to the stochastic steady state as soon as all sellers have reset their prices, as indicated by the dotted red line and the dashed green line. In combination, however, the two-month staggered pricing friction is amplified into 10 months of autocorrelated price changes, as shown in the solid blue line.

The gradual impulse response results from sellers only partially adjusting their list prices when they have the opportunity to do so in order to not ruin their chances of attracting a buyer by being substantially overpriced. Repeated partial adjustment results in serially correlated price changes that last far beyond the point that all sellers have reset their price.[46] Note that the impulse response

case and discuss alternative assumptions in Appendix A.6.

[46] With staggered pricing there are further dynamic incentives because price resetters leapfrog sellers with fixed

includes endogenous entry, which weakens house price momentum as buyers and sellers re-time their entry into the market to take advantage of the gradual price change. This effect, which is discussed in Section 1.6, is not strong enough to eliminate momentum. Appendix A.6 shows that concave demand also generates significant momentum for a downward shock. Intuitively, if other sellers do not lower their prices immediately, cutting a house's price substantially has a small effect on its probability of sale and leaves money on the table.

Despite the five-fold amplification, the staggered-pricing variant of the model only explains one quarter of the three-year impulse response in the data. This is unsurprising: there are many potential frictions that cause momentum, so it would be unrealistic to expect staggered pricing alone to be amplified by a factor of 20 in order to fully explain the data.[47]

To fully explain the impulse response in the data, I use the backward-looking variant of the model and raise the fraction of backward-looking sellers α until the impulse response function to the renter flow utility shock peaks after 36 months. This occurs when 26.5 percent of sellers are backward-looking. By contrast, without concave demand, between 78 and 93 percent of sellers would have to be backward-looking in order to explain a 36-month impulse response to the renter flow utility shock, with the precise number depending on how the non-concave demand curve is calibrated, as described in Appendix A.6.

Far fewer backward-looking sellers are needed to match the data with concave demand because the strategic complementarity creates a two-way feedback. When a shock occurs, the backward-looking sellers are not aware of it for several months, and the rational sellers only slightly increase their prices so that they do not dramatically reduce their chances of attracting a buyer. When the backward-looking sellers do observe increasing prices, they observe a much small increase than in the non-concave case and gradually adjust their price according to their AR(1) rule, reinforcing the incentives of the rational sellers not to raise their prices too quickly.

Panel B of Figure 10 compares the model with 26.5 percent backward-looking sellers in solid orange to the AR(5) impulse response in dotted black and a model with an identical fraction of

prices and are subsequently leapfrogged themselves. The interested reader is referred to Appendix A.4.8 for a detailed discussion of the dynamic intuition with staggered pricing.

[47] The momentum created by staggered pricing cannot be dramatically enhanced by increasing the length of staggering without increasing the average time to sale. For instance, if prices were fixed for four months instead of two months, the longer friction would be offset because there would be fewer sellers who remain stuck at an old price when a group of sellers sets their prices.

backward-looking sellers without concave demand in dashed turquoise. The impulse response with concave demand and the AR(5) impulse response are similar, although the model impulse response grows less at the beginning and is slightly more S-shaped than the AR(5) impulse response. This is the case because backward-looking sellers are insensitive to the shock for several months and so the growth rate of prices takes a few months to accelerate. Without concave demand, there is an immediate jump in prices as rational sellers raise their prices as soon as the shock to fundamentals occurs. This is followed by nine months of rapid price growth as the backward-looking sellers catch up. The strategic complementarity thus provides considerable amplification.[48]

With additional initial sources of price insensitivity it is likely that the 26.5 percent figure could be reduced even further.[49] Intuitively, concave demand creates an incentive to price close to others that interacts with *any* source of heterogenous price insensitivity to create additional momentum.[50] One particular friction that the literature has identified as causing momentum—incomplete information and learning by sellers and possibly buyers—merits additional discussion because the amplification from concave demand is likely to be particularly potent. In a model with dispersed information without strategic complementarities, such as Lucas' (1972) "islands" model, Bayesian learning about a change in fundamentals occurs fairly rapidly. Indeed, Anenberg (2013) shows that lagged market conditions do not have a significant impact on seller pricing after a four months. However, with a strategic complementarity and dispersed information, the motive to price close to others makes higher order beliefs—that is beliefs about the beliefs of others—matter, a point first made by Phelps (1983) and modeled by Woodford (2003) and Lorenzoni (2009). Learning about higher order beliefs is more gradual because agents must learn not only about fundamentals but also about what everyone else has learned. Strategic complementarities in such a framework can cause very gradual price adjustment even if first-order learning occurs rapidly.

[48] A direct empirical test of the degree to which concave demand amplifies momentum is beyond the scope of this paper. I do, however, have one intriguing data point that may point the way for such a test in the future: a smaller data set of merged Altos-DataQuick listings for Phoenix. Phoenix has a higher housing supply elasticity and by some measures exhibits less momentum than coastal California, and a preliminary analysis of its micro data suggests that the degree of curvature in Phoenix may be weaker. With many MSAs of data, one could evaluate whether the degree of concavity in a cross section of cities can explain differences in momentum across cities.

[49] For instance, adding staggered pricing to the rule-of-thumb model reduces the fraction of rule-of-thumb sellers needed to explain the data to 23.5 percent.

[50] Heterogeneity in sensitivity is key, as insensitivity that is uniform across identical sellers would imply that all sellers price at the average price, neutralizing the strategic complementarity.

1.6 Can Momentum Help Explain Housing Cycles?

This section argues that momentum helps explain the three striking features of the dynamics of housing cycles presented in Section 1.2: volume and inventory are more volatile than price, price changes and inventory levels are highly correlated, and inventory swings correspond to periods where seller entry and sales move in opposite directions. Momentum plays a role in causing these features because some buyers and sellers re-time their sales and purchases in light of predictable price changes. Before showing how this explains the three facts, I analyze the impulse response for non-price variables to provide intuition. The precise cause of momentum does not matter greatly for the re-timing of entry, so I focus on the backward-looking model with 26.5 percent backward-looking sellers because it fully captures the momentum in the data.

1.6.1 Impulse Responses of Non-Price Variables

Figure 11 shows the impulse responses of price, sales volume, months of supply, and buyer and seller entry in a frictionless model without backward-looking sellers (dotted red) and in a model with 26.5 percent backward-looking sellers (solid blue). Recall that the shock reduces the value of being a renter and increases the incentives to enter the market to buy.

Without staggered pricing, price jumps immediately and gradually returns towards the stochastic steady state, so there is not a strong incentive to buy or sell today relative to tomorrow. Buyer entry and seller entry, shown in panel D in dotted red and dash-dotted green, both jump on impact due to the change in the relative value of homeownership and the elevated house price. Buyer entry is slightly higher for 18 months as the ratio of buyers to sellers slowly rises until it settles on a stable transition path to the stochastic steady state. The slow adjustment of market tightness, in turn, causes a gradual increase in volume and decrease in months of supply.

By contrast, the momentum generated with a small fraction of backward-looking sellers makes price changes predictable. This creates a strong incentive for potential buyers on the margin of entering to enter today and for sellers on the margin of entering to wait to do so until prices rise.[51]

[51] In the model, this operates through the entry cutoffs c_t^* and k_t^*, which are defined by differences of value functions in equations (9) and (10). For instance, the cutoff cost for a renter to enter $k_t^* = V_t^r - V_t^b$. Because the value function of being a renter V_t^r accounts for the likelihood of getting a shock and entering as a buyer in the future, when prices are expected to rise V_t^r falls relative to V_t^b, k_t^* falls, and the mass of buyer entrants, which is proportional to $1 - K(k_t^*)$, rises.

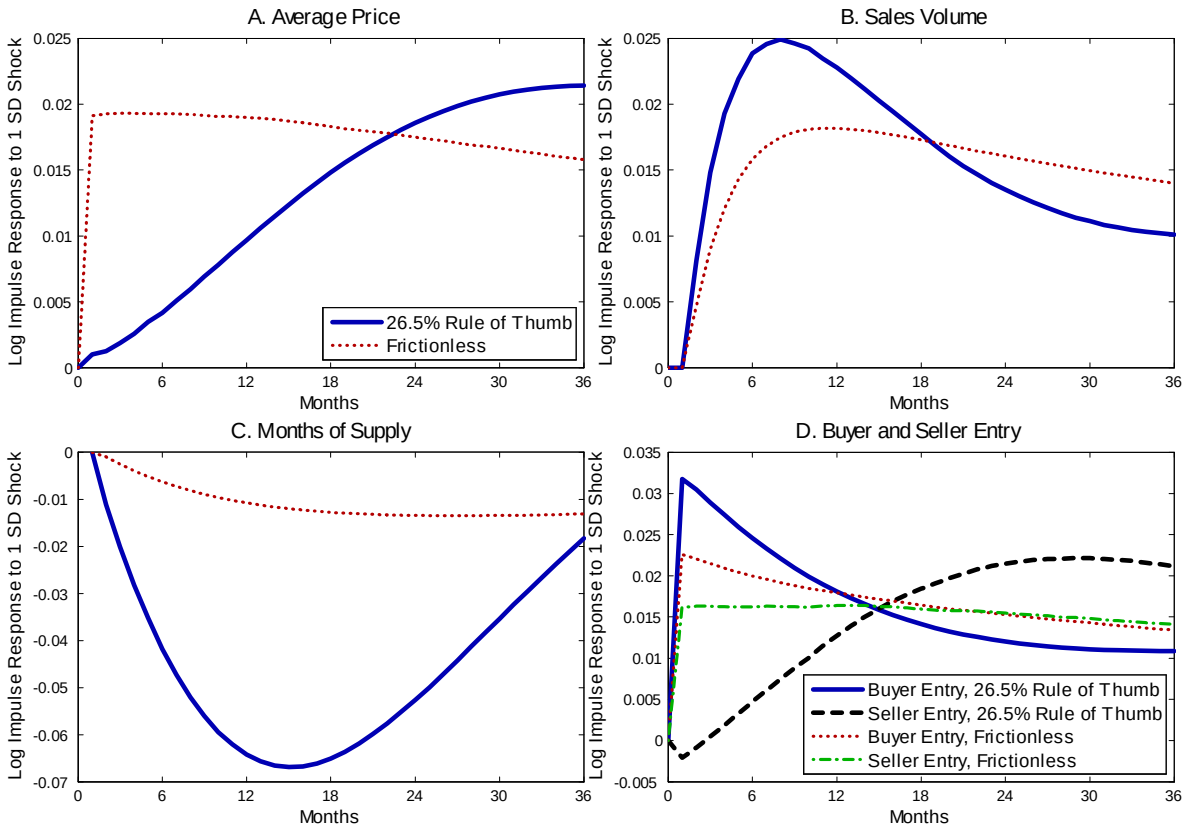

Figure 11: Impulse Response Functions in the Rule-ofThumb Model

Notes: Each panel plots the indicated impulse response to a one standard deviation shock for the frictionless and backward-looking variants of the model. The frictionless model uses the same calibration and shock as the 26.5 percent backward-looking model with no backward-looking sellers. Simulated impulse responses are calculated by differencing two simulations of the model from periods 100 to 150, both of which use identical random shocks except in period 101 in which a one standard deviation negative draw is added to the random sequence, and then computing the average difference over 100 simulations.

The entry responses are visible in panel D as a gap opens up between the solid blue line, which represents buyer entry with 26.5 percent backward-looking sellers, and the dotted red line, which corresponds to buyer entry in a frictionless model. A similar gap opens up for sellers, as shown by the dashed black line (26.5 percent backward-looking) and the dash-dotted green line (frictionless). Volume picks up and the growth in sales, overshooting of buyer entry, and undershooting of seller entry relative to the frictionless case together cause inventory to adjust more rapidly and substantially than it does in the frictionless model, as shown in panel C. The stock of renters becomes depleted and the stock of homeowners becomes enlarged to the extent that 15 months after the shock, they reverse roles and overshoot the frictionless price path again.

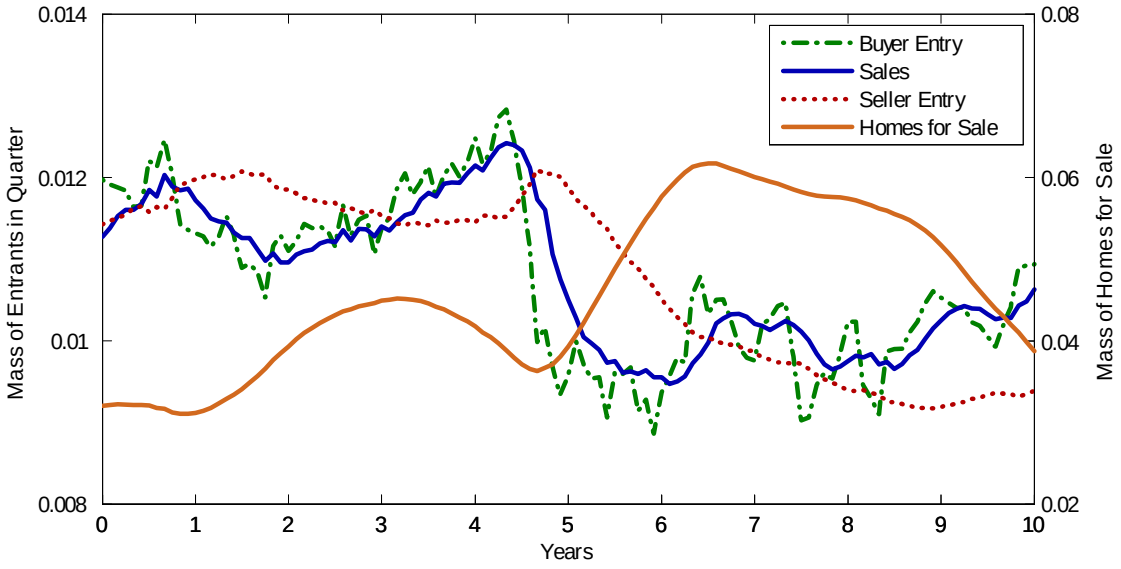

Figure 12: Sample Simulation of Entry, Volume, and Homes For Sale

Notes: This figure shows buyer entry, seller entry, sales, and the stock of homes listed for sale from a simulation of the backward-looking variant of the model with 26.5 percent backward-looking sellers.

This causes inventory and sales volume to mean revert more quickly. In fact, inventory and sales overshoot the frictionless impulse response once again as prices begin to stabilize due to the glut of sellers and lack of buyers. These responses look similar to the panel VAR presented in Section 1.2.

1.6.2 Explaining the Housing Cycle Facts

Buyer and Seller Entry Forward looking entry responses imply that seller entry and buyer entry move in opposite directions at peaks and troughs, corresponding to periods of sudden inventory adjustment, as shown for the recent boom and bust in Figure 2. While the impulse responses illustrate a similar pattern in the model, Figure 1.6.2 provides further confirmation of the model's ability to replicate the data by showing a sample simulation that looks strikingly similar to Figure 2. Initially there is a sellers' market in which inventory is low, and buyer and seller entry track one another. When the market peaks, buyer entry dries up but seller entry remains high as sellers seek to sell to buyers who need to buy now and are willing to pay high prices. As a result, inventory quickly spikes. When buyers finally re-enter the market, seller entry lags and the inventory glut dissipates.

Table 8: Quantitative Performance of Calibrated Models

	Data	Frictionless	Staggered	26.5% Backward Looking
SD Annual $\Delta \log$ (Real Price)	0.065	0.67	$0.065^{\diamondsuit}$	$0.065^{\diamondsuit}$
SD Annual $\Delta \log$ (Sales)	0.143	0.060	0.057	0.090
SD Annual $\Delta \log$ (Inventory)	0.207	0.040	0.055	0.306
Regression Coefficient of	-.140	0.124	0.010	-0.196
log (Inventory) on $\Delta_{yr} \log$ (Real Price)	(.015)	(.023)	(0.006)	(.007)
Regression R^2	.543	0.034	0.000	0.797

Notes: $^{\diamondsuit}$ indicates the model is calibrated to match the data. All are statistics calculated as means of 200 random simulations of 500 years. The standard deviations of annual log changes in the model are calculated by collapsing simulated data to the quarterly level, taking logs, and reporting the quarterly standard deviations of annual differences. Inventory is measured as months of supply. The regression of log price changes on log inventory levels is as in equation (26), with inventory measured as months of supply at the midpoint of the year differenced to calculate the log change in price. The frictionless model, which is not recalibrated to match the data, uses the staggered price calibration but would look nearly identical if the backward-looking calibration were used.

The Relative Volatility of Price, Volume and Inventory In housing search models without momentum, inventory and volume are too smooth relative to the data. The top half of Table 8 shows the standard deviation of annual changes of log price, log volume, and log months of supply in the data and three versions of the model. In the frictionless price model, months of supply is about a fifth as volatile and volume less than half as volatile as the data.

This is not unique to my particular frictionless model. In a broad class of housing search models, combining the steady-state value of being a seller with the steady-state price and differentiating yields:

$$\frac{dp}{d\Pr\left[Sell\right]} = \frac{Seller\ Surplus}{1-\beta}.$$

This steady state response illustrates that if the seller surplus is not miniscule, price is very sensitive to the probability of sale, which is mechanically related to inventory and volume.[52] The relative volatility of volume and inventory are also low due to the gradual dynamic adjustment of market tightness as shown in the impulse responses.

The low volatility of inventory is directly analogous to labor search models. Shimer (2005)

[52] Diaz and Jerez (2013) explain the relative volatilities in a model without momentum by using a calibration in which the seller surplus is 0.5% of the purchase price. Consequently, they argue that price is too insensitive and volume and time on the market are too sensitive to shocks and introduce a model with amplified price volatility. My calibration, which uses a seller surplus that is approximately 7.5% of the steady state price, implies that price is too volatile in a frictionless setting. Head et al. (2014) make a similar point that momentum reduces price volatility.

shows that unless the employer surplus is tiny, labor search models have difficulty accounting for the volatility of unemployment because most of the response to a shock is absorbed by the wage. Here, the unemployment rate is analogous to inventory and the wage rate is analogous to price.

Like sticky wages in labor search models, momentum makes house prices adjust more slowly and slightly reduces price volatility. Quantities adjust slightly more and inventory adjusts substantially more, as shown in the impulse responses and Table 8, which shows the standard deviation of annual changes for log price, log sales, and log inventory averaged over 200 500-year simulations. Inventory is somewhat too volatile in the calibrated model, although it is of the same order of magnitude as the data in contrast to the frictionless model. Volume, on the other hand, falls slightly short of the data, which suggests that other factors, such as lock in due to equity (Stein, 1995), may play a role in amplifying volume volatility. Appendix A.2 shows that the model's strongest prediction about relative volatilities—that momentum and inventory volatility are positively correlated—is borne out in the cross-section of cities used for the panel VAR in Section 1.2.

Housing Phillips Curve In the data, price changes are strongly negatively correlated with inventory levels, creating a "housing Phillips curve." In the frictionless case depicted in Figure 11, price changes are negatively correlated with inventory changes, albeit weakly because the inventory response is delayed due to gradual entry and search frictions. With persistent but mean reverting shocks, this generates a positive correlation between price changes and inventory levels because when inventory is high, prices are low and tend to rise towards the stochastic steady state. This can be seen in the bottom half of Table 8, which shows a regression coefficient β_1 in:

$$\Delta_{t,t-4} \log(p) = \beta_0 + \beta_1 \log(MS_{t-2}) + \varepsilon, \tag{26}$$

estimated on simulated quarterly data. For a frictionless model, the regression coefficient is significantly positive, although with a small R-squared.

With momentum, inventory rapidly adjusts and then mean reverts while price appreciation grows and then gradually weakens, as shown in Figure 11. This creates a strong negative correlation between price changes and inventory levels. Table 8 shows that with 26.5 percent backward-looking sellers, a robust negative relationship emerges. In fact, the relationship is slightly stronger than in

the data with an R-squared of nearly 0.8. Appendix A.2 shows that the housing Phillips curve is stronger, both in terms of the magnitude of β_1 and in terms of explanatory power, in cities with more momentum. This is consistent with the model.

1.7 Conclusion

The degree and persistence of autocorrelation in house price changes is one of the housing market's most distinctive features and greatest puzzles. This chapter introduces a mechanism that amplifies small frictions that have been discussed in the literature into substantial momentum. Search frictions and concave demand in relative price together imply that increasing one's list price above the market average is costly, while lowering one's list price below the market average has little benefit. This strategic complementarity induces sellers to set their list prices close to the market average. Consequently, modest initial price insensitivity to changes in fundamentals can lead to prolonged periods of autocorrelated price changes as sellers slowly adjust their list price to remain close to the mean.

This amplification mechanism depends critically on a concave effect of unilaterally changing a house's list price relative to the average on the probability of sale. I identify this effect in micro data by instrumenting for list price with a proxy for the equity position of sellers and find statistically and economically significant concavity.

To demonstrate the strategic complementarity's ability to prolong an initial source of price insensitivity, I introduce an equilibrium search model in which buyers avoid looking at homes they perceive to be overpriced. I calibrate the model to the micro data and consider the quantitative impact of two different sources of insensitivity. A two-month staggered pricing friction is amplified into ten months of autocorrelated price changes. If just 26.5 percent of sellers use a backward-looking rule of thumb, the impulse response to a shock lasts for three years, as in the data. Without concave demand, 78 to 93 percent of sellers would have to be backward-looking. The amplification channel also interacts with other frictions that have been discussed by the literature. In particular, concave demand in relative price would substantially amplify momentum created by learning in an "islands" model because learning about higher order beliefs is particularly sluggish. Assessing whether such a model can explain the momentum in the data without a small number of non-rational sellers is a promising path for future research.

Momentum has a substantial impact on housing dynamics because it causes forward-looking buyers and sellers to re-time their entry into the housing market in order to sell high and buy low. These buyer and seller entry patterns can help explain the relative volatilities of price, volume, and inventory, the "housing Phillips curve" relationship between price changes and inventory levels, and the sudden reversals between buyers' and sellers' markets that occurs at peaks and troughs.

Beyond the housing market, this chapter shows how a central idea in macroeconomics—that strategic complementarities can greatly amplify modest frictions—can be applied in new contexts. These contexts can, in turn, serve as empirical laboratories to study macroeconomic phenomena for which micro evidence has proven elusive. In particular, many models with real rigidities (Ball and Romer, 1990) use a concave demand curve. This chapter provides new evidence that a concave demand curve in relative price is not merely a theoretical construct and can have a significant effect on market dynamics.

Chapter 2:
How Do Foreclosures Exacerbate Housing Downturns?

2.1 Introduction

Foreclosures have been one of the dominant features of the recent housing market downturn. From 2006 through 2011, approximately 7.4 percent of the owner-occupied housing stock experienced a foreclosure.[53] Although the wave of foreclosures has subsided, foreclosures remain at elevated levels, and understanding the role of foreclosures in housing downturns remains an important part of reformulating housing policy going forward.

The behavior of the housing market concurrent with the wave of foreclosures is shown in Figure 2.1. Real Estate Owned (REO) sales – that is sales of foreclosed homes owned by banks and the GSEs – have made up between 20 and 30 percent of existing home sales nationally. Sales of existing homes fell 54.9 percent peak-to-trough; retail (non-foreclosure) volume fell 65.7 percent. Prices dropped considerably, with aggregate price indices plunging by a third and prices falling by a quarter for indices that exclude distressed sales. Time to sale and vacancy rates have also climbed, particularly in the retail market. Even with a slowdown in foreclosures due to lawsuits over fraudulent foreclosure practices, foreclosures have continued at a ferocious pace.

This chapter presents a model in which foreclosures have important general equilibrium effects that can explain much of the recent behavior of housing markets, particularly in the hardest-hit areas. By raising the number of sellers and reducing the number of buyers, by making buyers more choosey, and by changing the composition of houses that sell, foreclosures sales freeze up the market for retail sales and reduce both price and sales. Furthermore, the effects of foreclosures can be amplified considerably because price declines induce more default which creates further price declines, generating a feedback loop. A quantitative calibration suggests that these effects can be large: foreclosures exacerbate aggregate price declines by approximately 50 percent and retail price declines by 30 percent.

Despite the importance of foreclosures in the housing downturn, economists have not closely examined how the housing market equilibrates when there are a substantial number of distressed sales. A supply and demand framework, as employed by much of the financial literature on fire sales

[53] Data from CoreLogic. The data is described in Section 2.5 and Appendix B.4.

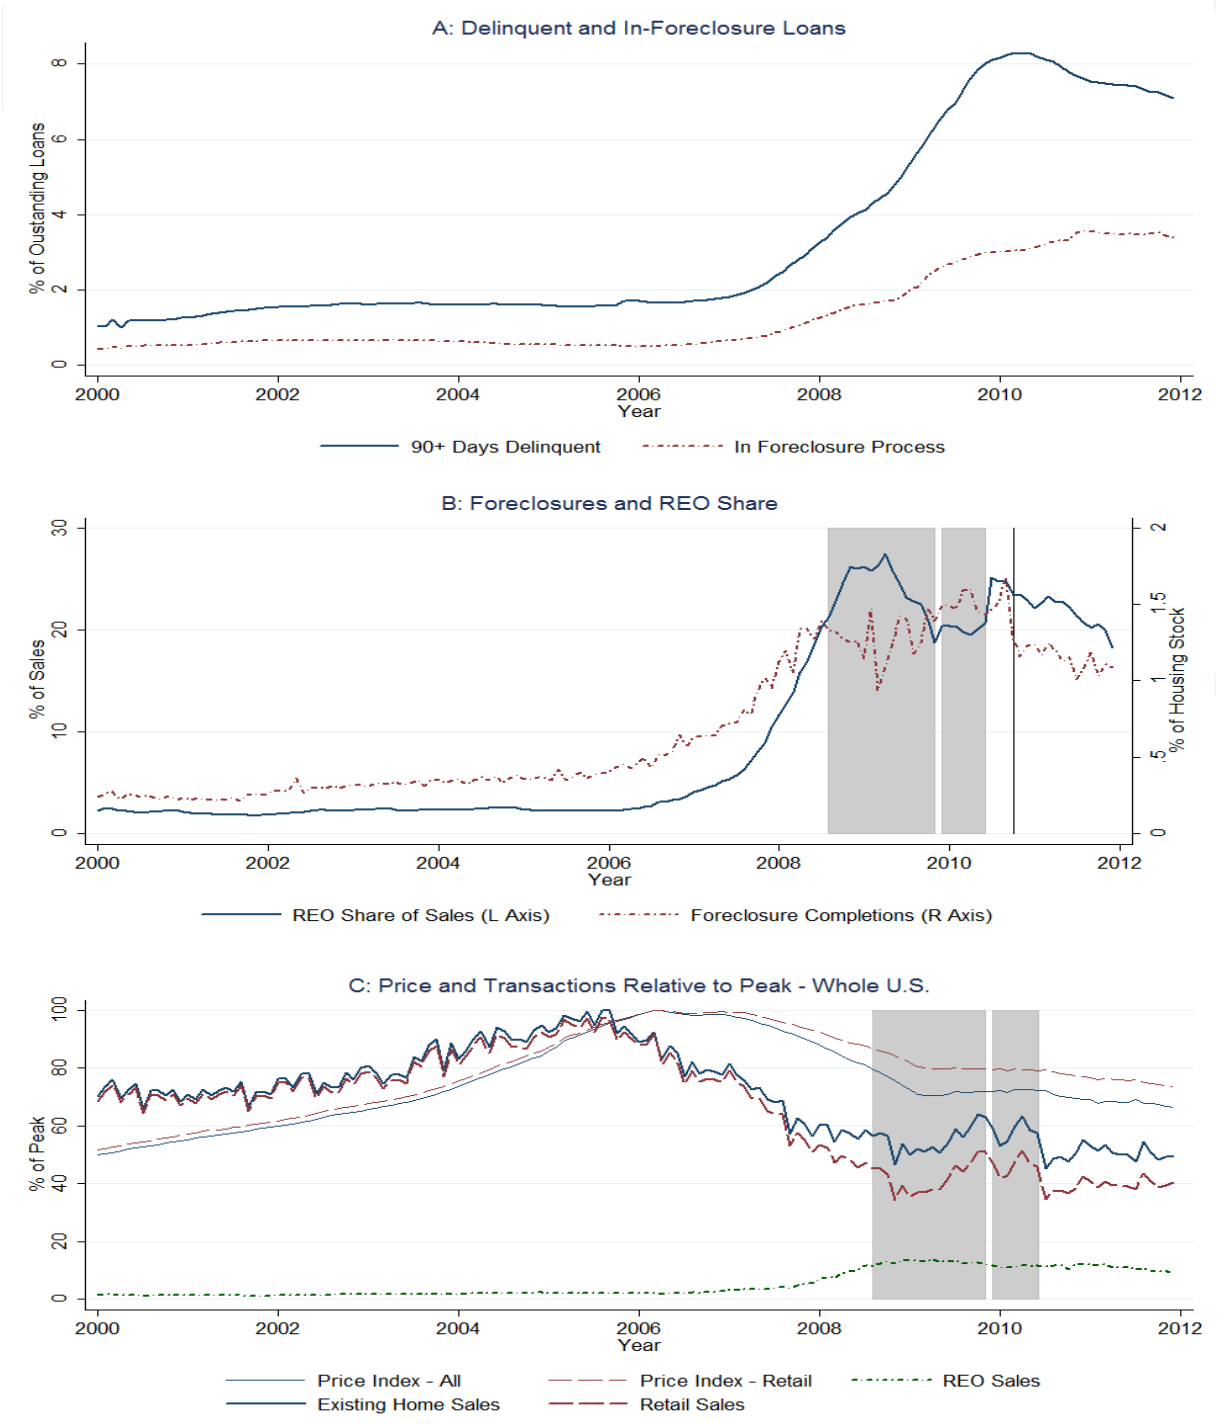

Figure 13: The Role of Foreclosures in the Housing Downturn

Notes: All data is seasonally adjusted national-level data from CoreLogic as described in the data appendix. The grey bars in panels B and C show the periods in which the new homebuyer tax credit applied. The black line in panel B shows when foreclosures were stalled due to the exposure of fraudulent foreclosure practices by mortgage servicers. In panel C, all sales counts are unsmoothed and normalized by the total number of existing home sales at peak while each price index is normalized by its separate peak value.

and illiquidity, can potentially explain declining prices and volumes with demand falling relative to supply but cannot speak to the freezing up of the retail market. Such models also assume that investors can adjust their positions continuously by transacting in a liquid market, yet housing is lumpy, illiquid, and expensive. A substantial literature has sought to adapt models to fit the peculiarities of the housing market and explain the positive correlation between volume and price. For instance, search frictions as in Wheaton (1990), Williams (1995), Krainer (2001), and Novy-Marx (2009), borrowing constraints as in Stein (1995), and nominal loss aversion as in Genesove and Mayer (2001) have been shown to play important roles in housing markets. Yet no paper has explicitly examined the role of distressed sales in a model tailored to housing.

To illustrate the mechanisms through which foreclosures affect the housing market, a simple model of the housing market with exogenous foreclosures is introduced. It adds two key ingredients to an otherwise-standard search-and-matching framework with stochastic moving shocks, random search, idiosyncratic house valuations, and Nash bargaining over price: REO sellers have higher holding costs and individuals who are foreclosed upon cannot immediately buy a new house. These two additions together dry up the market for normal sales, reduce volume and price, and imply that the market only gradually recovers from a wave of foreclosures. This occurs through three main effects. First, the presence of distressed sellers increases the outside option to transacting for buyers, who have an elevated probability of being matched with a distressed seller next period and consequently become more selective. This "choosey buyer effect" endogenizes the degree of substitutability between bank and retail sales. Second, because foreclosed individuals are locked out of the market, foreclosures reduce the likelihood that a seller will meet a buyer in the market through a "market tightness effect." This effect emphasizes that foreclosures do not simply add supply to the market: a key feature of foreclosures is that they also reduce demand. Third, there is a mechanical "compositional effect" as the average sale looks more like a distressed sale.

The choosey buyer effect in particular is novel and formalizes folk wisdom in housing markets that foreclosures empower buyers and cause them to wait for a particularly favorable transaction. For instance, The New York Times reported that "before the recession, people simply looked for a house to buy ... now they are on a quest for perfection at the perfect price," with one real estate agent adding that "this is the fallout from all the foreclosures: buyers think that anyone who is selling must be desperate. They walk in with the bravado of, 'The world's coming to an end, and I

want a perfect place.'"[54] The Wall Street Journal provides similar anecdotal evidence, writing that price declines "have left many sellers unable or unwilling to lower their prices. Meanwhile, buyers remain gun shy about agreeing to any purchase without getting a deep discount. That dynamic has fueled buyers' appetites for bank-owned foreclosures."[55] Although other papers such as Albrecht et al. (2007, 2013) and Duffie et al. (2007) have included seller heterogeneity in an asset market model, no paper that does so has generated a choosey buyer effect, which turns out to be important in explaining the disproportionate freezing up of the retail market.

To provide a more realistic treatment of the downturn, the basic model of the housing market is embedded in a richer model of mortgage default in which borrowers with negative equity may default on their mortgage or be locked into their current house despite a desire to move. This generates a new amplification channel: an initial shock that reduces prices puts some homeowners under water and triggers foreclosures, which cause more price declines and in turn further default. While reminiscent of the literature initiated by Kiyotaki and Moore (1997), the price declines here are caused by the general equilibrium effects of foreclosures. Lock-in of underwater homeowners also impacts market equilibrium by keeping potential buyers and sellers out of the market.

The richer model is used to quantitatively evaluate the extent to which foreclosures have exacerbated the ongoing housing bust. This quantitative analysis takes a two-pronged approach. First, we assess the strength of the amplification channel and its sensitivity to various parameters in the model. Second, we fit the model to data from the 100 largest MSAs to assess the empirical size of the amplification channel and test its implications across metropolitan areas. The model matches the data on the size of the price decline, the number of foreclosures, price declines in the retail market, and the REO share of sales. It also matches the heterogeneity in foreclosure discounts over the cycle found by Campbell et al. (2011). However, it falls short of explaining the full sales decline, suggesting that other forces have depressed transaction volume in the downturn. The quantitative analysis reveals that foreclosures exacerbate the aggregate price decline in the downturn by approximately 50 percent in the average MSA (or in other words account for a third of the decline) and exacerbate the price declines for retail sellers by over 30 percent.

Finally, we analyze the impact of the foreclosure crisis on welfare in our model and simulate

[54] "Housing Market Slows as Buyers Get Picky" June 16, 2010.
[55] "Buyer's Market? Stressed Sellers Say Not So Fast" April 25, 2011.

three foreclosure-mitigating policies: slowing down foreclosures, refinancing mortgages at lower interest rates, and reducing principal. While we do not conduct a full normative analysis, the simulations of these policies highlight the trade-offs faced by policy makers.

The remainder of the chapter is structured as follows. Before presenting the model, section 2.2 presents facts about the bust across metropolitan areas. To explain the data, the remainder of the chapter develops a model of how foreclosures affect the housing market, first focusing on mechanisms and then on magnitudes. Section 2.3 introduces a model of exogenous defaults, and section 2.4 explores the intuitions and qualitative implications of the model. In section 2.5, the basic model is embedded in a more complete model in which negative equity is a necessary condition for default, which creates a new amplification mechanism in the form of a price-default spiral. The chapter then turns to the magnitudes of the effects identified in sections 2.3-2.5. Section 2.6 calibrates the model and quantitatively analyzes the model's comparative statics and the strength of the price-default amplification channel. Section 2.7 takes the model to the national and cross-MSA data from the ongoing downturn. Section 2.8 considers welfare and foreclosure policy, and section 2.9 concludes.

2.2 Empirical Facts

The national aggregate time series of price, volume, foreclosure, and REO share presented in Figure 2.1 mask substantial heterogeneity across metropolitan areas in the severity of the housing bust and wave of foreclosures. To illustrate this, Figure 14 shows price and volume time series for four of the hardest-hit metropolitan areas. In Las Vegas, for instance, prices fell 61.5 percent, retail sales fell 84.0 percent, and the REO share was as high as 76.4 percent. Figure 14 also illustrates how foreclosure sales substitute for retail sales: retail sales rise as REO sales recede and fall as REO sales surge.

To provide more systematic facts about the heterogeneity of the bust across MSAs, we use a proprietary data set provided to us by CoreLogic supplemented by data from the United States Census. CoreLogic provides monthly data for 2000-2011 for the nation as a whole and the 100 largest MSAs, from which we drop 3 MSAs because the full data are not available for these locations at the start of the crisis. The data set includes a house price index,[56] a house price index for retail

[56] The CoreLogic price index is a widely-used repeat sales index that has behaved similarly to other cited indices

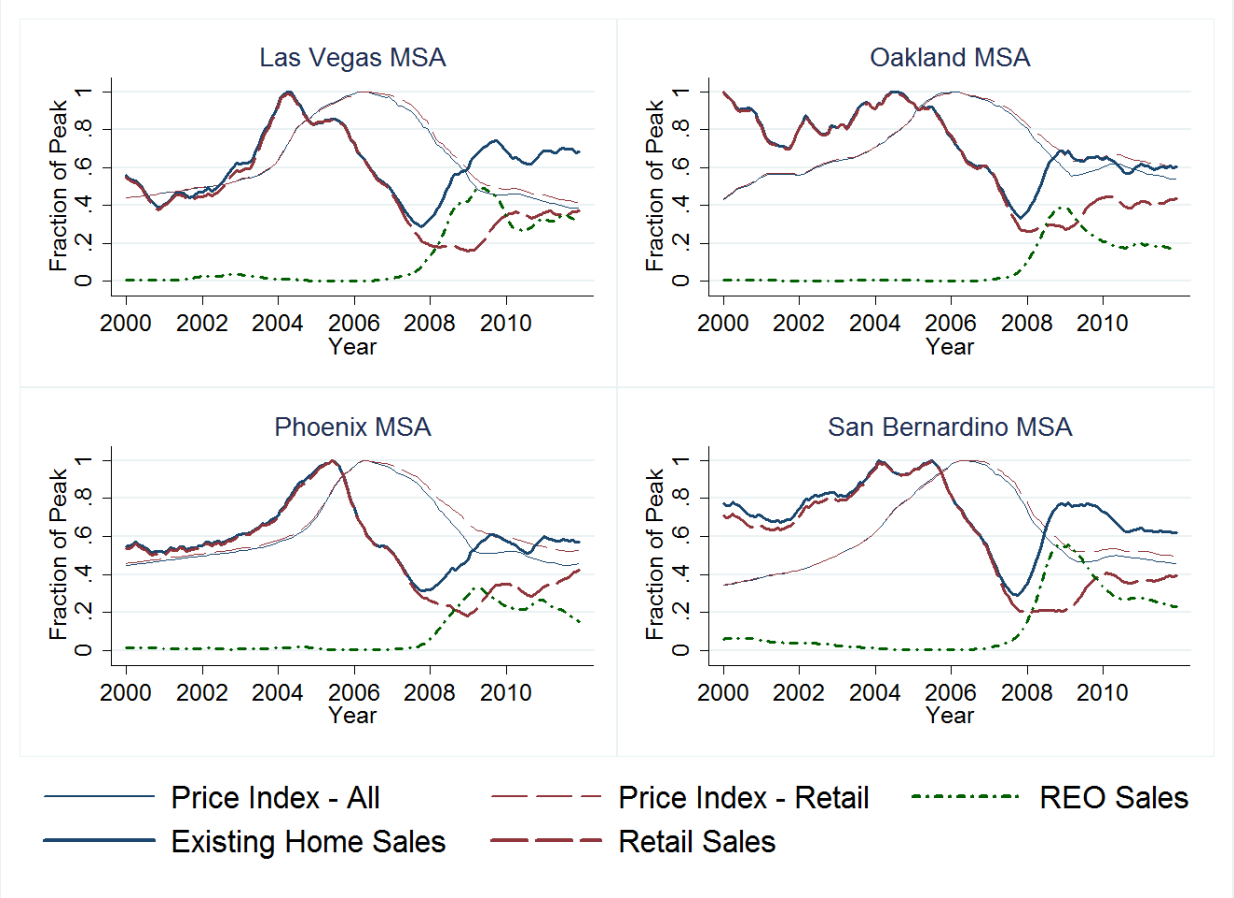

Figure 14: Price and Transaction Volume in Selected MSAs With High Levels of Foreclosures

Notes: All data is seasonally adjusted CBSA-level data from CoreLogic as described in the Appendix B.4. The sales lines are smoothed using a moving average. In panel C, all sales counts are normalized by the total number of existing home sales at peak while each price index is normalized by its separate peak value.

sales only,[57] the number of completed foreclosure auctions, sales counts for REOs, new houses, existing houses (including short sales), and the estimates of quantiles of the LTV distribution described previously. These statistics are compiled by CoreLogic using public records. CoreLogic's data covers over 85 percent of transactions nationally. We seasonally adjust the CoreLogic data and smooth the sales count series using a moving average. A complete description of the data is in appendix B.4.

By far the best predictor of the size of the bust was the size of the preceding boom. Figure 15

in that it fell by a third during the downturn. The S&P Case-Shiller index shows similar declines to the CoreLogic index. The FHFA expanded-data index, which includes FHFA data proprietary deeds data from other sources, fell 26.7 percent.

[57] Given the small number of distressed properties prior to the downturn, price indices for distressed properties are typically not estimated. The CoreLogic non-distressed price index drops REO sales and short sales from the database and re-estimates the price index using the same methodology.

plots the change in log price from 2003 to 2006 against the change in log price from each market's peak to its trough through 2011. There is a clear downward pattern, with the notable exception of a few outliers in the lower-left of the diagrams which correspond to metropoitan areas in southern Michigan which experienced a substantial bust without a large boom.

Figure 15 also reveals a more subtle fact in the data: places that had a larger boom had a more-than-proportionally larger bust. While a linear relationship between log boom size and log bust size has an r-squared of .44, adding a quadratic term that allows for larger busts in places with larger booms as illustrated in Figure 15 increases the r-squared to .57.

This chapter argues that by exacerbating the downturn in the hardest-hit places, foreclosures can explain much of why the why the relationship between log boom size and log bust size is not linear. This explanation implies an additional reduced-form cross-sectional test: because default is closely connected to negative equity, a larger bust should occur in locations with the *combination* of a large bubble and a large fraction of houses with high loan balances – and thus close to default – prior to the bust. To provide suggestive evidence that this prediction is borne out in the data, the points in Figure 15 are color-coded by quartiles of share of houses in the MSA with over 80 percent LTV in 2006. While the highest measured LTVs came in places that did not have a bust – home values were not inflated in 2006, so the denominator was lowest in these locations – one can see that the majority of MSAs substantially below the quadratic trend line were in the upper end of the LTV distribution.

To investigate whether the interaction of high LTV and a big bust combined is correlated with a deep downturn more formally, we estimate regressions of the form:

$$\begin{aligned} Y \quad = \quad & \beta_0 + \beta_1 \max \Delta_{03-06} \log(P) + \beta_2 \left[\Delta_{03-06} \log(P)\right]^2 \\ & + \beta_3 \left(\text{Z } \max \text{Share LTV} > 80\%\right) + \beta_4 \left(\Delta_{03-06} \log(\text{P}) \times \text{Z LTV} > 80\%\right) \\ & + \beta_5 \left(\text{Z \% Second Mortgage, 2006}\right) + \beta_6 \left(\Delta_{03-06} \log(\text{P}) \times \text{Z \% Second}\right) \\ & + \beta_7 \left(\text{Z Saiz Land Unavailability}\right) + \beta_8 \left(\text{Z Wharton Land Use Regulation}\right) + \varepsilon \end{aligned} \tag{27}$$

where Z represents a z-score and the outcome variable Y is either the maximum change in log price, the maximum change in log retail prices, the maximum change in log existing home sales, the maximum change in log retail sales, the maximum REO share, or the fraction of houses that

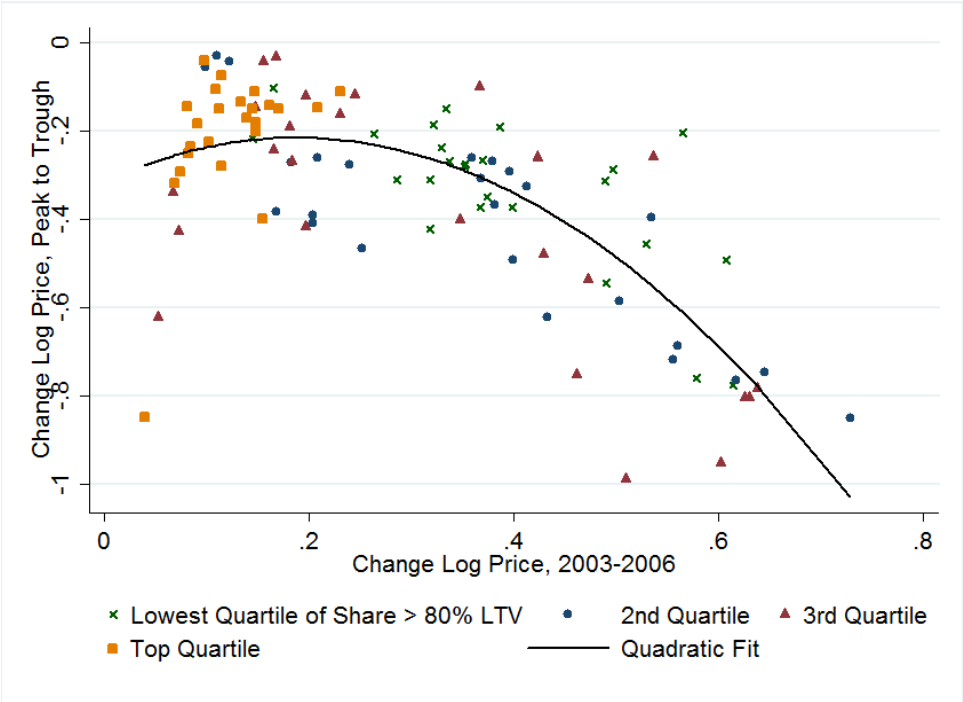

Figure 15: Boom vs. Bust Across MSAs

Note: Scatter plot of seasonally adjusted data from CoreLogic along with quadratic regression line. The data is fully described in Appendix B.4. Each data point is an MSA and is color coded to indicate in which quartile the MSA falls when MSAs are sorted by the share of homes with over 80% LTV in 2006. There are two main take-aways. First, there is a non-log-linear relationship between the size of the boom and the size of the bust, as the regression r-squared is .44 for a linear model and .57 for a quadratic model. Second, the color-coding of the points provides suggestive evidence that it is the combination of high LTV and a large bubble that is associated with a disproportionately large bust.

experience a foreclosure. The key coefficient is β_4. This regression is similar in spirit to Lamont and Stein (1999), who show that prices are more sensitive to income shocks in cities with a larger share of high LTV households, except rather than using income shocks to measure volatility, we use the size of the preceding bubble as measured by 2003-2006 price growth. We add the fraction of individuals with a second mortgage or home equity loan to the regression because these loans have received attention in analyses of the downturn (Mian and Sufi, 2011). Finally, to proxy for the housing supply elasticity we use a land unavailability index and the Wharton land use regulation index both from Saiz (2010). Table 9 shows summary statistics for our left hand side variables in the top panel and our right hand side variables in the bottom panel

The regression results are shown in Table 10. The first two columns show the impacts on price

Table 9: MSA Summary Statistics

	Unweighted Mean	SD	Min	Max	N
Max $\Delta \log(\text{P})$	-.3398355	.2292549	-.9895244	-.0286884	97
Max $\Delta \log(\text{P}_{\text{Retail}})$	-.2784659	.1929688	-.9229212	-.0388126	97
Max $\Delta \log(\text{Sales}_{\text{Existing}})$	-.9354168	.2684873	-2.53161	-.4671416	97
Max $\Delta \log(\text{Sales}_{\text{Retail}})$	-1.174493	.3181327	-2.736871	-.5613699	97
$\frac{\text{Sales}_{\text{REO}}}{\text{Sales}_{\text{Existing}}}$	.3147958	.1648681	.0834633	.795764	97
% Foreclosed	.0870826	.0719943	.0104154	.4205121	97
$\Delta \log(\text{Price})_{03-06}$	.2974835	.179294	.0389295	.7288995	97
Share LTV $> 80\%$	.1452959	.0756078	.025514	.3282766	97
Frac Second Mort, 06	.2026752	.0527425	.0259415	.2896224	97
Saiz Land Unav	.2779021	.2112399	.009317	.7964462	97
Wharton Land Reg	.2215807	.7050566	-1.239207	1.89206	97

Notes: Summary statistics for variables used in regression analysis. All data is from CoreLogic and fully described in Appendix B.4. Data is for 100 largest MSAs excluding three for which complete data are unavailable as described in the appendix.

and retail price. While the additional variables do not explain all of the non-log-linearity, they have substantial predictive power. The key coefficient shows that the interaction between a large bubble and the share of homes with high LTV is correlated with large price declines, as suggested by Figure 15. These interactions also have a large effect on the REO share of sales and fraction foreclosed, suggesting that foreclosures have something to do with these trends. The coefficient on land regulation is also negative yet small, reflecting the amplification provided by a high housing supply elasticity. Having many houses with a second mortgage also reduces prices.

For sales, the regression has noticeably less predictive power and the dominant term is the constant. As discussed in the analysis of the national calibration above, this suggests that foreclosures combined with the size of the bubble will do a much worse job explaining the volume decline than the price decline, something that will be borne out in our cross-MSA simulations. The interaction between LTV and the bubble is insignificant for existing sales but significant and negative for retail sales. The pattern of REO volume largely replacing retail volume is consistent with the four markets with high levels of foreclosure in Figure 14.

Table 10: Cross MSA Regressions on the Impact of the Size of the Bubble and Its Interaction With High LTV

Dependent Variable:	$\Delta \log(\text{P})$	$\Delta \log(\text{P}_{\text{Retail}})$	$\Delta \log(\text{Sales}_{\text{Existing}})$	$\Delta \log(\text{Sales}_{\text{Retail}})$	$\frac{\text{Sales}_{\text{REO}}}{\text{Sales}_{\text{Existing}}}$	% Foreclosed
$\Delta \log(\text{Price})_{03-06}$	1.501	0.884	-1.932	-0.493	-1.662	-0.615
	(0.656)**	(0.444)**	(0.450)***	(0.797)	(0.477)***	(0.189)***
$\Delta \log(\text{Price})^2_{03-06}$	-3.369	-2.440	1.431	-1.046	3.016	1.267
	(0.785)***	(0.529)***	(0.573)**	(0.936)	(0.561)***	(0.229)***
Z Share LTV $> 80\%$	0.062	0.064	-0.008	0.005	-0.009	-0.008
	(0.040)	(0.030)**	(0.038)	(0.054)	(0.033)	(0.015)
$\Delta \log(\text{P}) \times$ Z LTV $> 80\%$	-0.314	-0.314	-0.146	-0.379	0.218	0.197
	(0.155)**	(0.135)**	(0.138)	(0.182)**	(0.107)**	(0.067)***
Z Frac Second Mort, 06	-0.058	-0.037	-0.066	-0.071	-0.022	-0.019
	(0.033)*	(0.023)	(0.033)**	(0.044)	(0.027)	(0.010)*
$\Delta \log(\text{P}) \times$ Z Second	0.099	0.036	0.102	0.049	0.245	0.080
	(0.112)	(0.087)	(0.096)	(0.121)	(0.088)***	(0.036)**
Z Saiz Land Unav	-0.021	-0.008	-0.011	-0.014	0.003	-0.003
	(0.018)	(0.015)	(0.018)	(0.021)	(0.015)	(0.006)
Z Wharton Land Reg	-0.031	-0.027	0.017	-0.001	0.019	0.008
	(0.014)**	(0.011)**	(0.017)	(0.019)	(0.012)	(0.005)*
Constant	-0.417	-0.282	-0.551	-0.944	0.461	0.136
	(0.107)***	(0.073)***	(0.064)***	(0.132)***	(0.081)***	(0.031)***
r^2	0.644	0.703	0.298	0.353	0.521	0.606
N	97	97	97	97	97	97

Notes: * = 10% Significance, ** = 5% Significance *** = 1% significance. All standard errors are robust to heteroskedasticity. All data is from CoreLogic and fully described in Appendix B.4. Data is for 100 largest MSAs excluding 3 for which complete data are unavailable as described in appendix B.4. The key take-away is the coefficient on the interaction between bubble size and LTV, which shows that these two factors together were associated with a large price decline in the bust.

2.3 Housing Market Model

To theoretically examine the effect of forelcosures, we develop a model of the housing market in which foreclosures are exogenous. We subsequently embed this model in a framework in which default is modeled more realistically. Consequently, in this section, we focus on the mechanisms and qualitative predictions and defer a quantitative analysis of the model to section 2.7.

2.3.1 Setup

We consider a Diamond-Mortensen-Pissarides-style general equilibrium search model of the housing market. Search frictions play an important role in housing markets: houses are illiquid, most households own one house and move infrequently, buyers and sellers are largely atomistic, and search is costly, time consuming, and random. Additionally, the outside options of market participants are crucial in search models, so a search framework is well-suited to formalizing the choosey buyer effect described in the introduction.

Time is discrete and the discount factor is β. There are a unit mass of individuals and a unit mass of houses, both fixed. This is a good approximation of the the downturn, in which there has been a very low level of new construction and decreased migration.[58]

The setup of the model's steady state is illustrated schematically in Figure 16. Table 11 defines the model's key variables. To simplify the analysis, we assume no default in steady state, which is approximately the case when prices are stable or rising.[59] In steady state, mass l_0 of individuals are homeowners. Homeowners randomly experience shocks with probability γ that induce them to leave their house as in Krainer (2001) and Ngai and Tenreyro (2013).[60] We assume that these shocks occur at a constant rate and that only individuals who receive a moving shock search for houses. This assumption turns off the amplification channel identified by Novy-Marx (2009) through which endogenous entry and exit decisions by market participants in a search model create a feedback loop which magnifies the effects of fundamental shocks.

An individual who receives a moving shock enters the housing market as both a buyer with

[58] We do not consider the impact of long-run changes in the homeownership rate and retirement rate on the long-run equilibrium of the market, nor do we consider the long-run impact of new construction, both of which may be affected by the downturn and are important subjects for future research.

[59] Allowing default in steady state complicates the analysis but does not substantially change the results.

[60] Moving shocks are a reduced form for a number of different different life events that trigger a change in housing preference, such as the birth of children, death, job changes, and liquidity shocks.

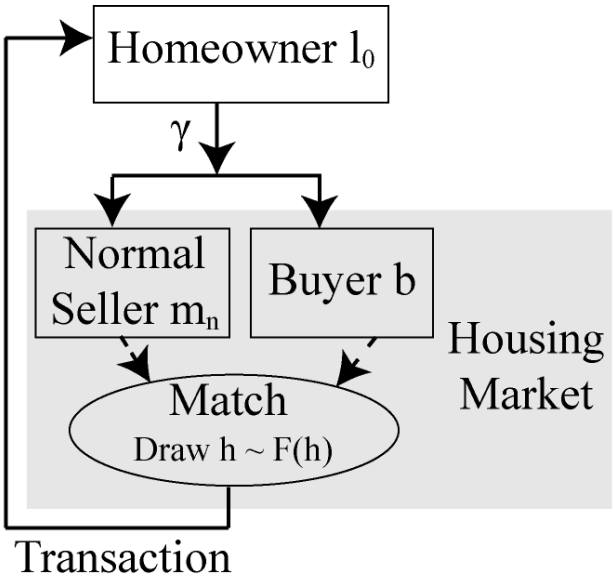

Figure 16: Schematic Diagram of No Foreclosure Steady State of Model

flow utility u_b and a normal seller with holding cost m_n. Because shocks create both a buyer and a seller, the model is a closed system with a fixed population.[61] In Section 2.7 we compare our model's predictions to data from both national and local markets, although the model as literally interpreted applies best to an metropolitan area.

As in Ngai and Tenreyro (2013), we assume that the buyer and seller act as independent agents. This means that there is no interaction between the buyer's problem or bargaining game and the seller's, and there is no structure placed on whether an individual buys or sells first. This assumption is not innocuous, as whether homeowners have sufficient liquidity to buy first may be important for market equilibrium and may affect bargaining as individuals who buy before selling holding two homes (Anenberg and Bayer, 2013). However, these effects are likely to be small relative to REOs. For instance, Springer (1996) examines several measures of seller motivation and finds that only REO sellers are distinguishable from normal sellers, and Anenberg and Bayer find that individuals who buy first sell their homes at a two percent discount, a figure that is swamped by the average REO discount.

Buyers and sellers in the housing market are matched each period. Matching is entirely random

[61] We use a closed system so that housing prices are not determined principally by the flow rates of buyers into and sellers out of the market but rather by the incentives of buyers and sellers in the market. Most moves are within-MSA (Sainai and Souleles, 2009) or to MSAs with highly correlated housing prices, so the assumption of a closed system is reasonable.

Table 11: Variables in Housing Market Model

Variable	Description
	Endogenous Variables
h	Stochastic Match Quality $\sim F(h)$
h_n, h_d	Cutoff h for normal, REO sellers
$S^B_{m,h}$, $S^S_{m,h}$	Surplus of type m seller with match quality h for buyer, seller
$p_{m,h}$	Price for type m seller with match quality h
μ	Market tightness (buyers/sellers)
$q_s(\mu)$, $q_b(\mu)$	Prob seller meets buyer, buyer meets seller
r_m, r_d	Ratio of normal, REO sellers to total sellers
l_0, l_1	Masses of homeowners, homeowners that could foreclose
v_b, v_n, v_d, v_r	Masses of buyers, normal sellers, REO sellers, renters
	Value Functions
V_h	Value of owning home with match quality h
V_n, V_d	Value of seller for normal, REO sellers
B	Value of buyer
R	Value of renter
	Parameters
β	Discount factor
γ	Probability of moving shock
α	Probability moving shock causes foreclosure
σ	Probability of leaving renting
θ	Seller's Nash bargaining weight
χ	Probability of match in period (C-D matching function)
ξ	Exponent in C-D matching function
λ	Parameter of exponential distribution for $F(h)$
a	Shifter on exponential distribution for $F(h)$
u_b, u_r	Flow utility of being a buyer, renter
m_n, m_d	Flow utility of being a seller for normal, REO

and search intensity is fixed, allowing us to focus on the effects of distressed sales rather than the search mechanism. When matched, the buyer draws a flow utility h from a distribution $F(h)$. Utility is linear and house valuations are purely idiosyncratic so that the transaction decision leads to a cutoff rule. These valuations are completely public, and prices are determined by generalized Nash bargaining. Because buyers know whether the seller is an individual or a bank in practice, symmetric information is reasonable. If the buyer and seller decide to transact, the seller leaves the market and the buyer becomes a homeowner in l_0 deriving flow utility h from the home until they receive a moving shock. If not, the buyer and seller each return to the market to be matched next period. Note that for simplicity we do not allow speculators or "flippers," who would presumably sell quickly.

We introduce foreclosures into this basic steady state setup by adding two key ingredients. First, REO sellers have a higher holding costs, which is the case for several reasons. Mortgage servicers, who execute the foreclosure and REO sale, have substantial balance sheet concerns. In most cases, they must make payments to security holders until a foreclosure liquidates, and they must also assume the costs of pursuing the foreclosure, securing, renovating, and maintaining the house, and selling the property (Theologides, 2010). Furthermore, even though they are paid additional fees to compensate for the costs of foreclosure and are repaid when the foreclosed property sells, the servicer's effective return is likely far lower than its opportunity cost of capital. Additionally, owner-occupants have much lower costs of maintenance and security. Finally, REO sellers usually leave a property vacant and thus forgo rental income or flow utility from the property.

An implicit assumption is that no deep-pocketed and patient market maker buys from distressed sellers and holds the property until a suitable buyer is found. While investors or "flippers" have bought some foreclosures, most have been sold by realtors to homeowners. This is likely due to agency problems and high transactions costs.

Second, individuals who experience a foreclosure are locked out of the market. This reflects the fact that a foreclosure dramatically reduces a borrower's credit score. Indeed, many banks, the GSEs, and the FHA will not lend to someone who recently defaulted. Instead, foreclosed individuals become renters. This is supported by the data: Molloy and Shan (Forthcoming) use credit report data to show that households that experience a foreclosure start are 55-65 percentage points less likely to have a mortgage two years after a foreclosure start. For simplicity, we assume that the rental market is segmented, and renters flow back into buying at an exogenous and fixed rate. While segmentation is a somewhat extreme assumption in the long run, it is a more reasonable approximation for the short-run effects in which we are primarily interested as conversions from owner occupied to rental units are costly and slow. Introducing an endogenous rental price and making the outflow rate covary with the price would create a force that mitigates some of the effects in the model.

Because flow utilities for foreclosures are drawn from the same distribution as for non-foreclosures, we are implicitly assuming that foreclosures are of roughly equal quality, which is likely not the case in practice (Gerardi et al., 2013). In our calibration, we are careful to use moments from studies that caerfully control for quality.

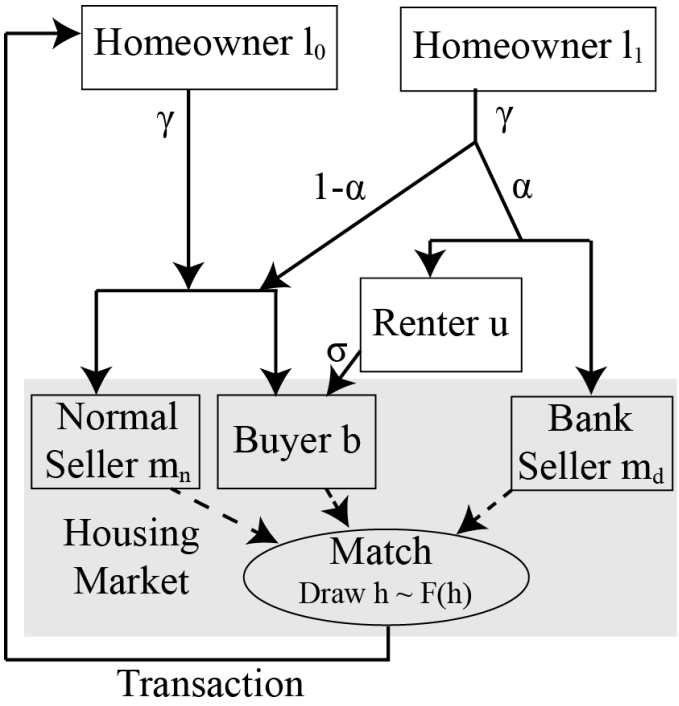

Figure 17: Schematic Diagram of Housing Market Model With Foreclosure

While these are the only two new assumptions we make, foreclosures may have other effects. They may cause negative externalities on neighboring properties due to physical damage, the presence of a vacant home, or crime. Campbell et al. (2011) show that such effects are small and highly localized, although contagion is certainly possible in neighborhoods with high densities of foreclosures. There may also be buyer heterogeneity with respect to their willingness to purchase a foreclosure, generating an additional channel through which the REO discount widens as non-natural buyers purchase foreclosures. Finally, foreclosures may cause banks to limit credit supply, as shown theoretically by Chatterjee and Eyigungor (2011).

The two critical assumptions are introduced into the model in Figure 17. To simplify the analysis, we assume away re-default.[62] Instead, we consider a mass of potential defaulters and analyze how these potential defaulters flow through the system. One can think of these potential defaulters as homeowners with high mortgage balances as will be the case in section 2.5. These individuals have mass l_1, and at time $t = 0$, when we introduce the exogenous foreclosure shock,

[62] We assume away re-default to keep the model consistent with the extended model in section 2.5. While this assumption does slightly increase the speed of convergence back to steady state over the course of the crisis, it does not substantively alter the quantiatative or qualitative results.

we move everyone in l_0 to l_1. Potential defaulters in l_1 also receive moving shocks with probability γ, but if they receive a moving shock it triggers a foreclosure with probability $\alpha(t)$ and is a normal moving shock with probability $1-\alpha(t)$. If it is a normal moving shock, the homeowners becomes a buyer and a seller as in steady state. A foreclosure shock, however, causes a bank or GSE with holding cost m_d to take possession of the house and enter the housing market and the homeowner to become a renter with flow utility u_r.[63] Renters become buyers each period with exogenous probability σ. Because there is no re-default, all buyers, including those who were formerly renters, are added to l_0 when they buy a house, so the model gradually returns to steady state.

Buyers and sellers of both types are matched in the housing market. Let $v_b(t)$, $v_r(t)$, $v_n(t)$, and $v_d(t)$ be the masses of buyers, renters, normal sellers, and REO sellers in the market at time t. Market tightness $\mu(t)$ is equal to the ratio of buyers to sellers:

$$\mu(t) = \frac{v_b(t)}{v_n(t) + v_d(t)}. \tag{28}$$

Unlike general equilibrium search models of the labor market in which market tightness is determined principally by a free entry condition for firms posting vacancies, here market tightness is determined by flows into renting due to default and out of renting at rate σ.

For the matching technology, we use a standard Cobb-Douglas matching function so that the number of matches when there are b buyers and s sellers is $\chi b^{\xi} s^{1-\xi}$. The probability a seller meets a buyer in a period with market tightness μ is given by $q_s(\mu) = \frac{\chi b^{\xi} s^{1-\xi}}{s} = \chi \mu^{\xi}$, and the probability a buyer meets a seller is $q_b(\mu) = \frac{\chi b^{\xi} s^{1-\xi}}{b} = \chi \mu^{\xi-1}$.

Let $V_h(t)$ be the value of being in a house with match quality h at time t, $V_m(t)$ be the value of being a seller of type m (either n or d) at time t, $B(t)$ be the value of being a buyer at time t, and $R(t)$ be the value of being a renter at time t. $V_h(t)$ is equal to the flow payoff plus the discounted expected continuation value:

$$V_h(t) = h + \beta\left\{\gamma\left[V_n(t+1) + B(t+1)\right] + (1-\gamma)V_h(t+1)\right\}. \tag{29}$$

[63] The bank must hold a foreclosure auction, but in the vast majority of cases the auction reserve is not met and the bank takes the house as an REO. For instance, Campbell et al. (2011) report that 82 percent of foreclosures in Boston are sold as REOs rather than at auction. For simplicity we assume all houses become REO.

The match surplus created when a buyer meets a seller of type $m = \{n, d\}$ and draws an idiosyncratic match quality of h at time t is a key value in the model. Denote this surplus by $S_{m,h}(t)$, the buyer's portion of the surplus by $S^B_{m,h}(t)$, and the seller's portion by $S^S_{m,h}(t)$. Let the price of the house sold if a transaction occurs be $p_{m,h}(t)$. The buyer's share of the surplus is equal to the value of being in the house minus the price and their outside option of staying in the market:

$$S^B_{m,h}(t) = V_h(t) - p_{m,h}(t) - u_b - \beta B(t+1). \tag{30}$$

The seller's share of the surplus is equal to the price minus their outside option of staying in the market:

$$S^S_{m,h}(t) = p_{m,h}(t) - m - \beta V_m(t+1). \tag{31}$$

Prices are set by generalized Nash bargaining with weight θ for the seller, so:

$$\frac{S^S_{m,h}(t)}{S^B_{m,h}(t)} = \frac{\theta}{1-\theta} \ \forall \ m. \tag{32}$$

Buyers and type m sellers will transact if the idiosyncratic match quality h is above a threshold value, corresponding to zero total surplus and denoted by $h_m(t)$. Because total surplus is:

$$S_{m,h}(t) = V_h(t) - (m + u_b) - (\beta B(t+1) + \beta V_m(t+1)) \tag{33}$$

the cutoff is implicitly defined by:

$$V_{h_m}(t) = m + u_b + \beta(B(t+1) + V_m(t+1)). \tag{34}$$

We can then define the remaining value functions. The value of being a type m seller is equal to the flow payoff plus the discounted continuation value plus the expected surplus of a transaction times the probability a transaction occurs. Because sellers meet buyers with probability $q_s(\mu(t))$ and transactions occur with probability $1 - F(h_m(t))$, V_m is defined by:

$$V_m(t) = m + \beta V_m(t+1) + q_s(\mu(t))(1 - F(h_m(t)))E\left[S^S_{m,h}(t) \,|h \geq h_m(t)\right]. \tag{35}$$

The most important aspect of V_m is that in a downturn $q_s(\mu)$ falls below its steady state value because foreclosures create renters rather than buyers ($\mu < 1$). The chance that a seller does not meet a buyer thus reduces the value of being a seller.

The value of being a buyer is defined similarly, although we must account for the fact that the buyer can be matched with two types of sellers. Let the probability of matching with a type m seller conditional on a match be $r_m(t) = \frac{v_m(t)}{v_n(t)+v_d(t)}$. B is defined by:

$$B(t) = u_b + \beta B(t+1) + q_b(\mu(t)) \sum_m r_m(t) (1 - F(h_m(t))) E\left[S^B_{m,h}(t) | h \geq h_m(t)\right]. \quad (36)$$

Because of random matching, as more REO sellers enter the market the weight on REO sellers in the buyer's value function r_d rises. REO sellers are be more likely to sell, so foreclosures raise the value of being a buyer. The decline in μ caused by foreclosures also raises $q_b(\mu)$, further increasing the value of being a buyer.

It is worth discussing what the implications of allowing buyers to direct their search towards foreclosures would be. A model with completely segmented REO and retail markets produces unreasonable parameter values. Intuitively, the REO and retail markets are linked by a buyer indifference condition that the probability of a match times the surplus must be the same in the REO market and the retail market. With a reasonable foreclosure discount, buyer indifference can only hold if the opportunity cost of waiting slightly longer for a distressed sale – the flow utility from being in that house – is implausibly high.

Furthermore, it is unlikely that any buyers look exclusively at one type of property. Instead, partially-directed search, in which buyers are able to direct their search to particular sub-markets in which the REO share of vacancies is higher than other sub-markets but is still not close to one, is most plausible. Examples of sub-markets include neighborhoods within a MSA or lower priced homes where there are likely to be more foreclosures. In this case, the effects we identify would be most pronounced in those sub-markets which had the highest REO share of vacancies, although there would be some spillovers because marginal individuals would switch to the REO-laden market. This is consistent with the findings of Landvoigt et al. (2012) that price declines in San Diego were stronger at the lower end of the market. We leave understanding the role of foreclosures for within-housing-market dynamics to future research.

The value of being a renter is defined as:

$$R(t) = u_r + \beta \left\{ \sigma B(t+1) + (1-\sigma) R(t+1) \right\}. \tag{37}$$

We will assume $u_r = u_b$, so that a renter is simply a buyer without the option to buy.

The conditional expectation of the surplus given that a transaction occurs appears repeatedly in the value functions. This quantity can be simplified as in Ngai and Tenreyro (2013) by using (29) together with (33):

$$S_{m,h}(t) = V_h(t) - V_{h_m}(t) = \frac{h - h_m(t)}{1 - \beta(1-\gamma)}.$$

The conditional expectation is

$$E\left[S_{m,h}(t) \,|h \geq h_m(t)\right] = \frac{E\left[h - h_m(t) \,|h \geq h_m(t)\right]}{1 - \beta(1-\gamma)}. \tag{38}$$

We parameterize $F(\cdot) \sim \exp(\lambda) + a$, an exponential distribution with parameter λ shifted over by a constant a. The memoryless property of the exponential distribution implies that $E\left[S_{m,h}(t) \,|h \geq h_m^*(t)\right] = \frac{1}{\lambda}$. This is a fairly strong assumption. By using the exponential distribution in our simulations, we eliminate changes in the expected surplus due to changes in tail conditional expectations of the F distribution, which cannot be observed.

The model is completed with the laws of motion for the mass of sellers of type m, buyers, renters, and homeowners of type l_i. These laws of motion, which formalize Figure 17, are in appendix B.1.2.

Prices can be backed out by using Nash bargaining along with the definitions of the surpluses and (38) to get:

$$p_{m,h}(t) = \frac{\theta(h - h_m(t))}{1 - \beta(1-\gamma)} + m + \beta V_m(t+1) \tag{39}$$

This pricing equation is intuitive. The first term contains $h - h_m(t)$, which is a sufficient statistic for the surplus generated by the match as shown by Shimer and Werning (2007). As θ increases, more of the total surplus is appropriated to the seller in the form of a higher price. This must be normalized by $1 - \beta(1-\gamma)$, the effective discount rate of a homeowner. The final two terms

represent the value of being a seller next period, which is the seller's outside option. These terms form the minimum price at which a sale can occur, so that all heterogeneity in prices comes from the distribution of h above the cutoff $h_m(t)$. Because with the exponential distribution $E[h - h_m(t)] = \frac{1}{\lambda}$, all movements in average prices work through $V_m(t+1)$.

2.3.2 Numerical Methods

For reasonable parameter values, the model has a unique steady state that can be solved block recursively and studied analytically. The full derivation and existence and uniqueness proofs for the steady state can be found in appendix B.1.1. Although there are no foreclosures in steady state, the price and probability of sale for a REO seller are well defined and represent what would occur if a measure zero mass of normal sellers were instead REO sellers. For a fixed idiosyncratic valuation h, REO properties sell faster and at a discount due to the higher holding costs of distressed sellers.

The dynamics of the model, however, have no analytic solution, so we turn to numerical simulations. We solve the model using Newton's method as described in appendix B.1.2.

Simulating the model requires choosing parameters. We defer a more rigorous quantitative analysis to section 2.7, which features a richer model, and focus on the mechanisms at work in this section. Consequently, for now we present simulation results using an illustrative calibration similar to the one described in section 2.5. We simulate a wave of foreclosures by moving everyone in l_0 to l_1 at time $t = 0$ and raising α for a period of five years. After the wave of foreclosures, the model returns to the original steady state.

2.4 Basic Model Results and Mechanisms

2.4.1 Market Tightness, Choosey Buyer, and Compositional Effects

The qualitative results in our model are caused by the interaction of three different effects: the "market tightness effect," the "choosey buyer effect," and the "compositional effect." Each is crucial to understand the effect of foreclosures on the housing market.

First, because foreclosed individuals are locked out of the housing market as renters and only gradually flow back into being buyers, foreclosures reduce market tightness $\mu(t)$. This mechanically decreases the probability a seller meets a buyer in a given period and triggers endogenous responses

as each party's outside options to the transaction changes, altering the bargaining and the hs for which a sale occurs. For sellers, the reduction in market tightness reduces the value of being a seller for both types of seller, reducing prices and causing sellers to sell more frequently. The endogenous response is stronger for REO sellers who have a higher opportunity cost of not meeting a buyer. For buyers, the elevated probability of meeting a seller raises their expected value, leading to lower prices and a shift in the cutoffs that makes buyers more choosey.[64] The market tightness effect elucidates that an important element of foreclosures is a reduction in demand relative to supply, as in a typical market a move creates a buyer and a seller while foreclosures create an immeiate bank seller but a buyer only when the foreclosed upon individual's credit improve. This contrasts with some market analysts who treat foreclosures as a shift out in supply rather than a reduction in today's demand.

Second, the value of being a buyer rises because the buyer's outside option to transacting, which is walking away and resampling from the distribution of sellers next period, is improved by the prospect of finding an REO seller who will give a particularly good deal. Mathematically, as REOs make up a larger fraction of total vacancies, r_d rises and the term in the sum in (36) relating to REO sales gets a larger weight. This term is larger because REO sellers are more likely to transact both in and out of steady state. The resulting increase in buyers' outside options leads buyers to become more aggressive and demand a lower price from sellers in order to be willing to transact. In equilibrium, this leads to buyers walking away from more sales. Importantly, this effect will be most prevalent in the retail market where sellers are less desperate and therefore less willing to accommodate buyers' demand for lower prices, resulting in a freezing up of the retail market.

The choosey buyer effect is new to the literature. Albrecht et al. (2007, 2013) introduce motivated sellers into a search model, but focus on steady-state matching patterns (eg whether a high type buyer can match with a low type seller) and asymmetric information regarding seller type. Duffie et al. (2007) consider a liquidity shock similar to our foreclosure shock, but a transaction occurs whenever an illiquid owner meets a liquid buyer, and so while there are market tightness effects their model does not have a choosey buyer effect.

[64] In the calibration utilized here and in later sections, we set $\xi = .84$. Thus the effect on $q_s(\mu)$ significantly outweighs the effect of market tightness on $q_b(\mu)$.

The market tightness effect and choosey buyer effect are mutually reinforcing. As discussed above, the market tightness effect is more pronounced for REO sellers. Because the value of being an REO seller falls by more, REO sellers become even more likely to sell relative to non-REO sellers during the downturn. This sweetens the prospect of being matched with an REO seller next period, amplifying the choosey buyer effect.

Finally, a greater share of REO sales makes the average sale look more like REO properties, which sell faster and at lower prices both in and out of steady state. Foreclosures thus cause a mechanical compositional effect that affects sales-weighted averages such as total sales and the aggregate price index.

The market tightness effect is the aspect of the model that comes closest to a standard Walrasian analysis with a single market for housing. By reducing the number of buyers relative to sellers, it is similar to an inward shift in the demand curve relative to the supply curve that reduces both prices and transaction volume. The market tightness effect does, however, asymmetrically impact REO and retail sellers due to their differential holding costs, leading to a greater freezing up of the retail market as buyers walk away from retail sellers in hopes of contacting increasingly-desperate REO sellers. These types of differential effects and further feedback loops – which stem from the choosey buyer effect and its interaction with the market tightness effect – are novel to the literature and differentiate our model from a simpler Walrasian model.

Furthermore, all three effects dissipate more slowly than in traditional asset pricing models because they depend on flows as well as stocks and lead to a sluggish return of the housing market to steady state. The choosey buyer and compositional effects last as long as foreclosures remain in the market, which is only a few months after the shock ends as these houses sell quickly. However, the market tightness effect persists for much longer as it takes several years for the renters to return to being homeowners.

2.4.2 Qualitative Results

Figure 18 shows the effect of a the five-year wave of foreclosures. Because the model is entirely forward looking, prices and probability of sale conditional on a match fall discretely on the impact of the shock at $t = 0$. This is typical in completely forward-looking models. The sluggish adjustment of house prices to shocks remains a puzzle for much of the literature, and a solution to this problem

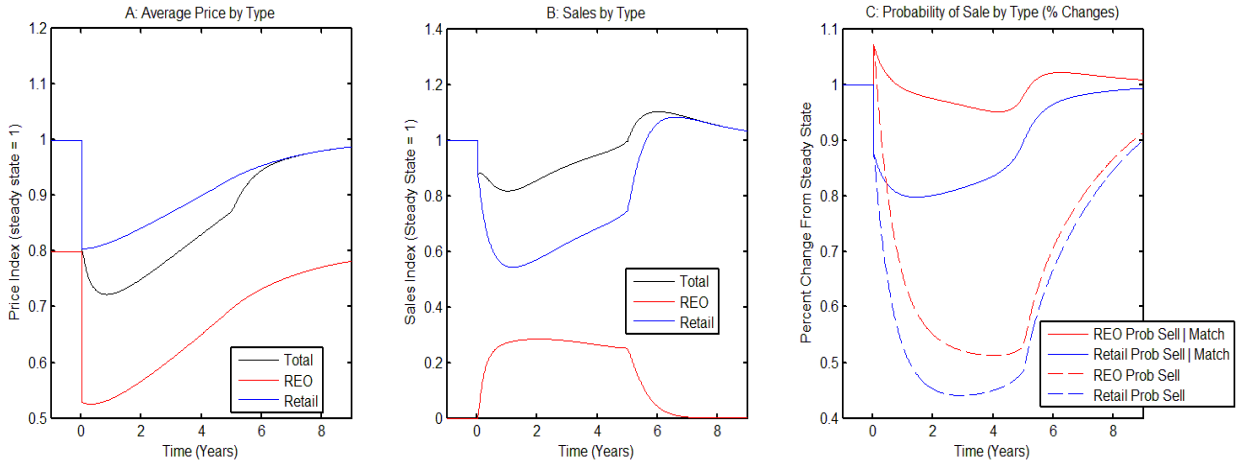

Figure 18: Housing Market Model: Qualitative Results

Notes: This figure shows the results of the housing market model with exogenous foreclosures using an illustrative calibration similar to the one developed in Section 2.5 and a five-year foreclosure shock. Panels A and B show the average price and sales by type, with pre-downturn price and volume normalized to 1. Prices drop discretely at time zero as is standard in forward-looking models with no uncertainty. The REO discount widens, the aggregate price index is pulled towards the REO index as REOs make up a greater share of the market, and prices rise in anticipation of the end of the downturn. Retail volume plunges dramatically, but the decline is partially made up for by surging REO volume. Panel C shows the probability of sale conditional on a match and the unconditional probability of sale for each type with the pre-downturn probability normalized to one. This panel illustrates the mechanisms at work in the model, as described in the main text. The key take-away is that the probability of sale conditional on a match, which is the clearest indicator of how the behavioral responses of buyers and sellers play out in equilibrium, falls dramatically for retail and is roughly flat for REO.

is outside the scope of this chapter.

As shown in panel A, at $t = 0$ prices fall considerably for both REO and retail and gradually return to steady state over the next several years. The overall sales-weighted price index dips more than retail sales as foreclosures are averaged in. The price movements lead to a substantial rise in the average REO discount that falls off over time.

Prices fall due to all three effects. Recall that from (39), movements in the average price of properties sold by a type m seller are controlled by movements in $V_m(t+1)$. The market tightness effect has a direct effect on the value of being a seller and thus brings down prices. Because this effect is stronger for REO sellers, this contributes to the larger REO discount. The choosey buyer effect has an indirect effect on $V_m(t+1)$, as in general equilibrium increased buyer choosiness reduces the value of being a type m seller, which causes prices to fall. The effect of market tightness on the value of being a buyer operates in a similar manner. Finally, there is a pure compositional effect as REO sales become a greater share of total sales, which is shown graphically

by the departure of the aggregate price index from the price index for retail sales.

As for sales, the wave of foreclosures sales causes the retail market to freeze up, with retail volume falling substantially as shown in panel B and REOs constituting a larger fraction of total sales than of total vacancies. Total volume, however, does not fall as much because much of the decline in retail sales is offset by REO sales. After the foreclosures end, sales return back to normal in a matter of months as REOs are eliminated from the market. Most of the sluggish adjustment comes from the dissipation of the accumulated renters and retail sellers, which takes several years.

The intuition behind the effects on transaction volume is more nuanced as the market tightness, choosey buyer, and composition effects have cross-cutting impacts. Panel C, which shows percent changes from steady state in the probability of sale both raw and conditional on a match, elucidates the role of each effect.[65]

Consider first the probability of sale conditional on a match, controlled by $h_m(t)$.[66] The market tightness effect on the probability a seller meets a buyer raises the probability of sale conditional on a match because sellers meet buyers less frequently and thus have a greater incentive to sell when they are matched, an effect which is stronger for REO sellers. The choosey buyer effect and the effect of market tightness on the probability a buyer meets a seller both reduce the probability of sale conditional on a match as buyers become more choosy. Panel C shows that the two effects offset for REO sales as the probability of sale conditional on a match fluctuates around its steady state value, while the choosey buyer effect and the market tightness effect on buyers dominate for retail sales as the probability of sale conditional on a match falls substantially. The relative strength of these two effects for the two types of sellers thus plays an important role in freezing up the retail market.

The market tightness effect, however, plays an additional role: it mechanically reduces volume because there are fewer buyers. This causes the unconditional probability of sale and thus transaction volume to fall for both types, although it falls more for REO sellers. Note, however, that decline for retail sales is quicker and the trough lasts longer.

The compositional effect also plays an important role in determining transaction volume. Because REOs sell faster both in and out of steady state, as the average sale looks more like an REO,

[65] The probability of sale conditional on a match is $\exp(-\lambda(h_m(t)-a))$ and the total probability of sale is $q_s(\mu(t))\exp(-\lambda(h_m(t)-a))$

[66] Time to sale is inversely related to the unconditional probability of sale.

volume rises. This is the main reason why total volume does not fall so dramatically. It is possible for volume to rise, although for reasonable calibrations we find that the market tightness effect is strong enough relative to the compositional effect that REO sales do not make up the full shortfall in retail sales and overall volume falls.

Qualitatively, the model explains many salient features of the housing downturn. The substantial decline in both retail and REO prices is consistent with the data in Figure 2.1, and the widened distressed sale discount in a downturn is corroborated by Campbell et al. (2011). The freezing up of the retail market and the large share of REO sales in total sales relative to listings is borne out in the data, as are a rise in times to sale and increasing vacancy rates. The fact that REO sales replace a good deal of the lost volume in the retail market is consistent with the evidence from the hardest hit markets as shown in Figure 14.

2.4.3 Isolating the Role of Each Effect

To further illustrate how each effect contributes to our results, Figure 19 depicts simulations identical to our main results for a wave of foreclosures except with the market tightness effect, choosey buyer effect, and both the choosey buyer and market tightness effects shut down. Although the market tightness effect plays an outsized role, all three effects are necessary for our results.

The market tightness effect generates a significant fraction of the price and volume declines. Row two also illustrates that the market tightness effect increases the conditional probability of sale for REO sellers during the downturn. Market tightness effects also cause total volume to decline because of the mechanical decrease in matching probabilities.

However, the choosey buyer effect plays an essential role in freezing up the retail market. As can be seen from row two, with no choosey buyer effect the conditional probability of sale for retail sellers essentially remains flat. On the other hand, from row one we can see that when only the choosey buyer effect is present there is a non-trivial decrease in this conditional probability. This freezing up is even more pronounced when both market tightness and choosey buyer effects are present due to their interaction.

The compositional effect mainly reduces the aggregate price index, as shown in row 3 of Figure 19. It also increases total volume slightly because REO sales sell faster.

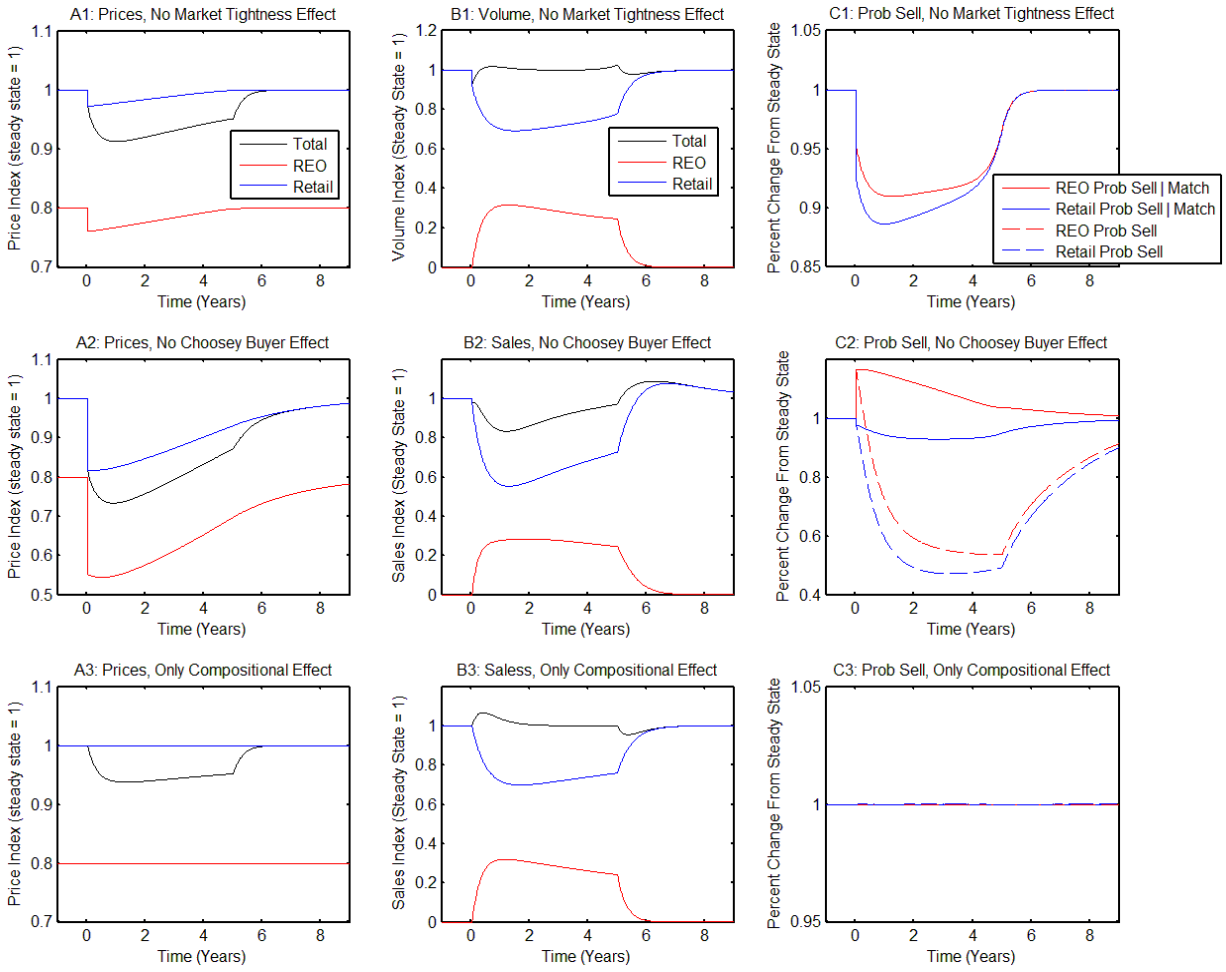

Figure 19: Isolating the Role of Each Effect

Notes: The top row shows price, sales, and conditional and unconditional probability of sale, all normalized to 1 for their pre-downturn values, for the case of no market tightness effect. The second row shows the same results for no choosey buyer effect. The third row shuts down both, leaving only the compositional effect. The figure shows that both the market tightness and outside option effects are critical for the qualitative results; in particular panel C2 shows that without the choosey buyer effect the probability of sale conditional on a match does not fall much for retail sellers. To shut down the market tightness effect, instead of creating renters we instead assume that distressed sale shocks create REO sellers and home-buyers. To shut down the choosey buyer effect, we modify the buyer's value function so that agents behave as if the probability they will hit a distressed seller is zero regardless of the presence of distressed sellers in the market. Because we calibrate to a steady state with no distressed sales, the steady state of these modified models replicates the steady state of our full model. See appendix B.2.1 for full details on these models.

2.5 An Extended Model of Default

Foreclosures are not random events. With few exceptions, negative equity is a necessary but not sufficient condition for foreclosure (Foote et al., 2008). This is because a homeowner with positive equity can sell his or her house, pay off the mortgage balance, and have cash left over without having to default. Homeowners with negative equity, however, are not able to pay the bank and thus default if they experience with a liquidity shock.

The previous section showed that foreclosures have general equilibrium effects that cause prices – and thus homeowner equity – to fall. In a world in which negative equity leads to foreclosure, this will cause more foreclosures and price declines, generating a feedback loop that amplifies the effects of an initial decline in house prices.[67]

In this section we embed the housing market component of the exogenous default model developed in the section 2.3 into a model in which negative equity is a necessary but not sufficient condition for default. Subsequent sections provide a rigorous quantitative analysis of the extended model and analyze welfare and foreclosure policy using the model.

2.5.1 Default in the Extended Model

We model default as resulting primarily from shocks that cause homeowners with negative equity to be unable to afford their mortgage payments, the so-called "double trigger" model of mortgage default. While "ruthless" or "strategic default" by borrowers has occurred, much of the literature on default argues that strategic default has contributed surprisingly little to foreclosures, particularly at low levels of negative equity.[68] Bhutta et al. (2010) use a method of controlling for income shocks to estimate that the median non-prime borrower does not strategically default until their equity falls to negative 67 percent. Even among non-prime borrowers in Arizona, California, Florida, and Nevada who purchased homes with 100 percent financing at the height of the bubble – 80 percent of whom defaulted within 3 years – over 80 percent of the defaults were caused by income shocks. Similarly, Foote et al. (2008) show that in the Massachusetts housing downturn

[67] In all cases we have considered, each additional round of feedback is smaller than the previous one generating a convergent series and a unique dynamic equilibrium, although in principle the feedback could be strong enough to generate a divergent series.

[68] Relevant papers than analyze the default decision and conclude that a "ruthless exercise" option model of default is insufficient include Deng et al. (2000), Bajari et al. (2009), Elul et al. (2010), and Campbell and Cocco (2014).

of the early 1990s, the vast majority of individuals who default have negative equity but most individuals with negative equity do not default. Consequently, the largest estimate of the share of defaults that are strategic is 15 to 20 percent.[69] To keep the model tractable, we thus do not model strategic default, nor do we model the strategic decision of the bank to foreclose or short sales.[70]

Modeling negative equity requires that homeowners have loan balances. We assume that homeowners in l_1 have a distribution of loan balances L defined by a CDF $G(L)$.[71] So that no foreclosures occur without an additional shock, in general we assume that $G(L)$ has continuous support on $[0, V_n^*]$, where the steady state value of being a normal seller which is equal to the expected price net of the costs of sale. We assume away re-default so that we do not need to worry how new home purchases affect $G(L)$.

To incorporate liquidity shocks into our model, we assume that they occur to individuals with negative equity at Poisson rate γ_I. All other shocks are taste shocks that occur at Poisson rate γ, so that liquidity shocks are in addition to normal shocks.

Liquidity and taste shocks have different effects depending on the equity position of the homeowner. Homeowners with any shock with $L \leq V_n(t)$ have positive equity enter the housing market as a buyer and seller. To keep the model tractable, we assume that buyers and sellers are identical once they pay off their loan balance. Homeowners with $L > V_n(t)$ have negative equity net of moving costs and default if they experience an income shock because they cannot pay their mortgage or sell their house. Defaulters enter the foreclosure process.[72] Although foreclosure is not immediate and some loans in the foreclosure process do "cure" before they are foreclosed upon, for simplicity we assume that foreclosure occurs immediately. We alter this and introduce foreclosure backlogs in Section 2.8. We also assume that income shocks are a surprise, so an underwater homeowner

[69] This estimate comes from Experian-Oliver Wyman. Guiso et al. (2013) analyze a survey that asks people whether they strategically defaulted and find that 25 to 35 percent of defaults are strategic.

[70] "Fishing" – that is listing a home for a high price and hoping that someone who overpays for it will come along as in Stein (1995) – and short sales are unusual because they require sellers to find a buyer who will pay a minimum price, which affects bargaining. Modeling short sales and their effect on market equilibrium is an important topic for future research.

[71] We are agnostic as to the source of the loan balance distribution and leave this unmodeled. $G(L)$ is fixed over time because principal is paid down slowly, particularly by those in the upper tail of the loan balance distribution who are relevant for the size of the feedback loop.

[72] It is also possible for banks to possess the house and rent it to the homeowner or to offer a short sale, in which the bank accepts a sale at a price below the oustanding loan balance. While these options have become more popular in recent years due in large part to political pressure, for most of the crisis the banks simply foreclosed on borrowers.

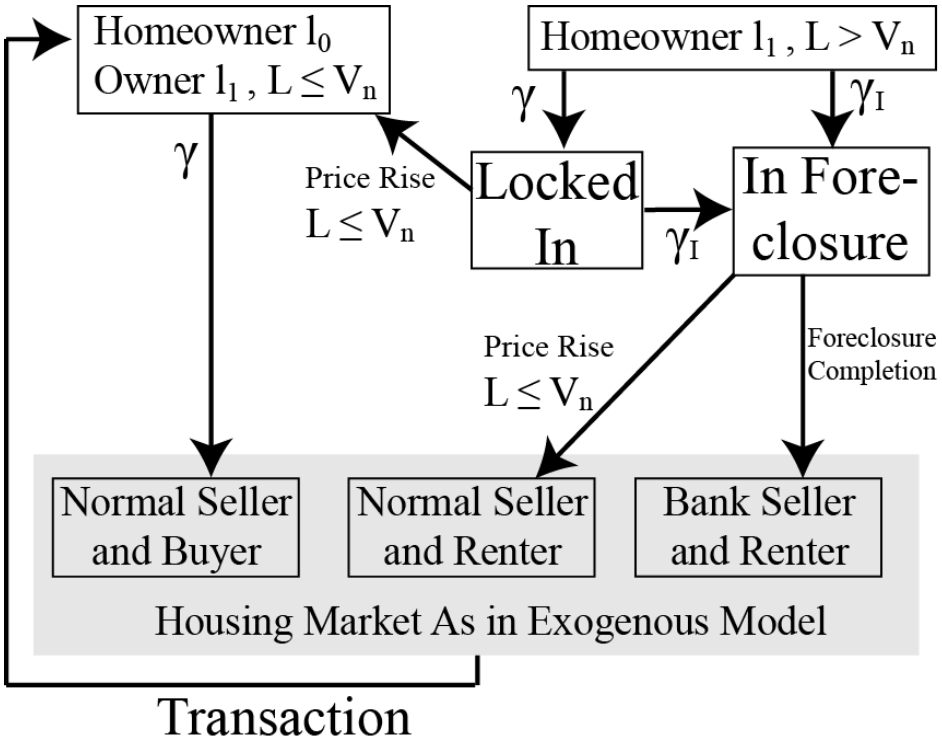

Figure 20: Extended Model Schematic Diagram

expecting an income shock cannot list their house with the hope of getting a high-enough price that they can pay off their loan before the bank forecloses. While this may happen infrequenlty, it is unlikely very unlikely that a desperate seller would receive such a high price.

Finally, homeowners with negative equity who receive a taste shock would like to move but owe more than their house is worth. Consequently, they are "locked in" their current house until prices to rise to the point that they have positive equity.[73] We assume that once they do not move when they get a taste shock, these homeowners make accommodations and thus do not immediately move when they reach positive equity. Instead, once they have positive equity they become indistinguishable from households in l_0 who are waiting for a moving shock. In our context, lock-in further reduces sales in markets with many foreclosures.

Figure 20 shows our formalization of default in the extended model in a schematic diagram. The extended model only alters the mechanisms through which homeowners default and enter the housing market, and consequently the housing market component of the model is exactly as in

[73] Formally, define $w(t)$ as the mass of individuals who are locked in at time t. The distribution of loan balances in $w(t)$ will be the same as $G(L)$ truncated below at $V_n(t)$.

Table 12: Variables Used In Extended Model

Variable	Description
	Endogenous Variables
L	Loan Balance $\sim G(L)$
w	Mass of locked in homeowners
f	Mass of homeowners in foreclosure process
	Parameters
γ_I	Probability of a liquidity shock
δ	Probability above-water homeowner becomes renter
b_a, b_b	Parameters of Beta distribution for $G(L)$

section 2.3 and is thus not depicted in Figure 20. This structure preserves all of the key intuitions developed in section 2.4. Because there is still no default in steady state, the steady state remains the same. Because the modified model has an identical housing market to the model in section 2.3, the Bellman equations and cutoff conditions are unchanged. The new laws of motion are in Appendix B.1.4, and the additional parameters introduced in the extended model are listed in Table 12.

2.5.2 Starting the Downturn

An exogenous shock is required to generate an initial price drop. We introduce the exogenous shock in two different ways.

First, we assume that due to both tighter lending practices and income shocks a fraction $\delta(t)$ of individuals who sell their house as a normal seller after receiving a taste shock cannot buy a house and instead transition from owning to renting. This generates an initial market tightness effect that reduces prices, putting some individuals underwater and triggering a price-default spiral. Such a shock fits most naturally in the model. We use a 5-year increase in $\delta(t)$ to perform a sensitivity analysis in section 2.6.

However, when we take the model to the data in section 2.7, it is clear that the main shock in the recent episode is a bursting housing bubble. Our primary interest is understanding the amount of overshooting of prices caused by the presence of foreclosures and not realistically modeling the bursting of the bubble absent foreclosures. Consequently, we introduce a bursting bubble into the model with a permanent decline in a, the minimum flow utility of housing. While the flow utility of housing clearly did not fall overnight, the source of the initial drop in prices is immaterial to

our results and reducing a is the simplest way to generate such a price drop. The reduction in a should thus be seen as a stand-in for a number of factors that have depressed home prices, from changes in credit availability to belief disagreements to irrational exuberance.[74]

With a permanent decline in a and a constant hazard of an income shock, all individuals whose loan balance is above the new steady-state price level will eventually default. This does not seem realistic – many homeowners will eventually pay down their mortgage and avoid default. Rather than modeling the dynamic deformation of the $G(\cdot)$ distribution over time, we instead assume that after 5 years the hazard of income shocks γ_I gradually subsides over the course of a year.[75]

With both exogenous shocks, defaults due to negative equity and the resulting foreclosures amplify the effects of shocks in the housing market due to a price-foreclosure feedback loop. Due to the forward-looking nature of agents in the model, this spiral is capitalized into the prices of retail sales and REO sales when the shock occurs, with further gradual declines as the REO share of volume in the market increases.

2.6 Quantitative Analysis: Calibration and Amplification

Having analyzed the mechanisms at work in our extended model, we now turn to the model's implied magnitudes by conducting a quantitative analysis. In this section, we calibrate the model and examine the strength of the amplification channel. In section 2.7, we take the model to the data on the ongoing downturn.

2.6.1 Calibration

In order to simulate the model numerically, we must choose parameters. As mentioned previously, we parameterize the distribution of idiosyncratic valuations $F(\cdot)$ as an exponential distribution with parameter λ shifted by a. We parameterize the loan balance distribution $G(L)$ as a beta distribution with parameters b_a and b_b scaled to have support on $[0, V_n^*]$. This flexible distribution allows us to approximate the loan balance distribution in various locations on the eve of the crisis as described below. This gives 12 exogenous parameters to calibrate for the basic housing market

[74] While a bursting bubble is the most likely source of a large price change that would put many homeowners underwater, the type of foreclosure crisis we describe could be created by any type of large negative price shock.

[75] Formally, after 5 years γ_I falls by 5% of its previous value every month, taking roughly a year to settle at zero.

model – a, γ, m_d, m_n, u_r, u_b, θ, β, σ, λ, χ, and ξ – and three parameters to calibrate unique to the extended model – γ_I, and b_a and b_b. We also must choose the initial shock.

Our calibration procedure proceeds in three steps. We take care to calibrate to pre-downturn moments whenever possible in order to make our tests out-of sample, although in some cases we have no choice but to choose a parameter using data from the housing bust. First, we set γ, u_r, u_b, β, χ, ξ, and γ_I independently to match several moments. Second, we choose a, m_n, m_d, θ and λ so that the steady state of the model matches additional targets. Third, we calibrate b_a and b_b to the appropriate geographic unit that we are considering.

This leaves two variables that we do not choose through calibration: σ, the probability of leaving the rental market, and the initial shock. Although there are several guidelines regarding how long banks deny mortgages to individuals who default,[76] there is no good data on this parameter in practice. Consequently, we pursue a two-pronged approach. First, to understand the impact of σ and the exogenous shock, we examine the response of the model to different values of each. Second, we use data on the size of the bust across housing markets to select a preferred calibration of these two parameters, as described in Section 2.7. In the remainder of this section, we describe the moments to which we calibrate the model in our three steps.

Step 1: Exogenous Parameter Choices

We choose u_r, u_b, β, and γ so that one period is equivalent to one week, although the results are insensitive to the time interval chosen. We set the annual discount rate equal to 5%, so that the discount factor is $\beta = 1 - \frac{.05}{52}$. We assume $u_r = u_b = 0$ so that buyers and renters are identical in their flow utility. Buyers, however, have the option to buy which has considerable value so $B > R$. This assumption is equivalent to assuming that buyers and renters pay their flow utility in rent in a perfectly competitive rental market.

We set the probability of moving houses in a given week to fit a median tenure for owner occupants of approximately 9 years from the American Housing Survey from 1997 to 2005, so $\gamma = \frac{.08}{52}$.

We set $\xi = .84$ using estimates from Genesove and Han (2012), who use National Association of Realtors surveys to estimate the contact elasticity for sellers with respect to the buyer-to-seller

[76] Three years of good credit are needed to get a Federal Housing Administration loan, and according to Fannie Mae guidelines issued in 2010, individuals are excluded from getting a mortgage for two to seven years if they are foreclosed upon, depending on the circumstances. However, these guidelines are not ubiquitous.

ratio. χ is then a constant chosen to make sure the probability of matching never goes above 1 for either side of the market. We choose $\chi = .5$, which in our simulations leads to matching probabilities on $[0, 1]$. The results are robust to alternate choices of χ.

The one parameter that we need data from the downturn to choose is γ_I. We set $\gamma_I = 8.6\%$ annually using national data from CoreLogic on the incidence of foreclosure starts for houses with negative equity as described in appendix B.4.

Step 2: Matching the Pre-Downturn Steady State

We then fit the following five aggregate moments from the housing market prior to the housing bust to the model's steady state to set θ, a, λ, m_n, and m_d:

1. The mean house price for a retail sale, which we set equal \$300,000 as an approximation to Adelino et al.'s (2012) mean house price of \$298,000 for 10 MSAs. In reporting results, we normalize this initial house price to 1. Our results are approximately invariant to the mean house price as long as the residual variance is rescaled proportionally.

2. The residual variance of house prices due to the idiosyncratic preferences of buyers. We set this equal to \$10,000.

3. The REO discount in terms of mean prices, which we set equal to a quality-adjusted 20% based on the average discount in good times from Campbell et al. (2011), whose results are consistent with a literature surveyed by Clauretie and Daneshvary (2009). In Section 2.6 we also consider a 10% discount, closer to the estimates of Zillow (2010) and Clauretie and Daneshvary (2009).

4. Time on the market for retail houses, which we set to 26 weeks as in Piazzesi and Schneider (2009). This number is a bit higher than some papers that use Multiple Listing Service Data such as Anenberg (2013) and Springer (1996), likely because of imperfect adjustment for withdrawn listings and re-listings. Our results are not sensitive to this number.

5. Time on market for REO houses, which we set to 15 weeks. Springer (1996) analyzes various forms of "seller motivation" such as relocation and financial distress using data form Texas from 1991-3. He finds that a foreclosure sales are the only motivated sellers that have significantly reduced time on the market. His estimate is that time on the market is reduced

Table 13: Calibrated Parameter Values

a	γ	m_d	m_n	u_r	u_b	θ
3.399	$\frac{0.08}{52}$	-0.403	-0.091	0	0	0.087
β	λ	χ	ξ	γ_I	b_a	b_b
$1-\frac{0.05}{52}$	3.490	0.5	0.84	$\frac{.086}{52}$	0.898	1.223

by .2135 log points or 23.7%. However, REO sales are almost never withdrawn from the market, whereas retail sales are frequently withdrawn (Anenberg, 2013). We also attempt to adjust so our number excludes extremely rundown properties that sit on the market for several years.

These moments provide a unique mapping to θ, a, λ, m_n, and m_d, as described in appendix B.1.3.

The calibration procedure results in the parameter values listed in Table 13, with all prices and dollar amounts in thousands of dollars. Two main things are of note about the calibration. First, $m_d < m_n$, so the flow cost of being a REO seller is higher than the flow cost of being a regular seller. This is due to the fact that distressed sales take less time to sell and trade at a discount in steady state. Second, θ is quite low in order to rationalize the 20 percent discount for REO sales in steady state. This means the buyer will get a majority of the surplus and the value of being a buyer in the buyer's market of the downturn will be high.[77]

Step 3: Geographically-Specific Parameters

To calibrate the two parameters of the loan balance distribution b_a and b_b at the national and MSA level we use proprietary data from CoreLogic on seven quantiles of the combined loan-to-value distribution for active mortgages in 2006. These LTV estimates are compiled by CoreLogic using public records, with the LTV estimates supplemented using CoreLogic's valuation models.[78]

[77] A low θ is consistent with the logic of directed search and Genesove and Han's (2012) estimate of $\xi = .84$. In directed search models θ is determined endogenously through price posting behavior. The buyer's bargaining power is the elasticity of contacts with respect to market tightness ξ and so θ should be below $1 - \xi = .16$. Intuitively the seller has very low bargaining power because the seller's contact rate is more sensitive to a change in market tightness than the buyer's.

[78] Although the CoreLogic estimates of negative equity and the loan-to-value distribution are most cited, some have argued that they do not fully capture the extent of negative eqiuity. Recent estimates by Zillow use credit report data instead of public records and get a higher figure for negative equity than CoreLogic (methodological differences are described here http://blogs.wsj.com/developments/2012/05/24/negative-equity-more-widespread-than-previously-thought-report-says/). Beyond issues of data sourcing, Korteweg and Sorensen (2013) argue that traditional methods of estimating house price indices under-estimate the variance of the house price distribution and thus under-estimate the number of loans with high LTV.

Table 14: Sensitivity Analysis: Time Out of Market For Renters

	Price Index Decline	Total Sales Decline
$\boldsymbol{\sigma = 1/52}$ (1 year out)	4.0%	4.4%
$\boldsymbol{\sigma = 1/65}$ (1.25 years out)	4.9%	4.8%
$\boldsymbol{\sigma = 1/78}$ (1.5 years out)	5.8%	5.2%
$\boldsymbol{\sigma = 1/91}$ (1.75 years out)	6.6%	5.5%

Because our model concerns the entire owner-occupied housing stock and not just houses with an active mortgage, we supplement the CoreLogic data with the Census' estimates of the fraction of owner-occupied houses with a mortgage from the 2005-2007 American Community Surveys. We construct the empirical CDF of the loan balance distribution and then fit a beta distribution with parameters b_a and b_b to the empirical distribution using a minimum distance method described in Appendix B.4. The fit is quite good across the 50th to 100th percentiles of the LTV distribution.] Table 13 shows the resulting b_a and b_b for the nationwide loan balance distribution.

2.6.2 Strength of Amplification Channel and Comparative Statics

Having calibrated the model, we now gauge the potential magnitude of the amplification channel and elucidate key comparative statics in the extended model. Our initial shock will be one in which a fraction $\delta(t)$ of individuals who receive a taste shock transition from owning a home to renting, as described previously. This causes an initial price decline since it reduces the number of buyers but not the number of sellers. In particular, we use a shock of $\delta(t) = .10$ for five years.[79] Holding the initial shock constant, we then vary the average weeks out of the market for a renter, the steady state discount on REO sales, and the loan balance distribution. Specifically, we consider average times out of the market of 1, 1.25, 1.5, and 1.75 years and REO discounts of 10% and 20%. For the loan balance distribution, we consider beta distributions fitted to match the national data as well as data from New York, which had a low share of high LTV homes, and Las Vegas, which had a high share.

Since our initial shock is a market tightness effect, the strength of this effect will be governed by σ. Table 14 reports the price index decline and total volume decline generated by the initial shock in the absence of any defaults.

[79] Raising the size of the initial shock within reasonable parameter ranges magnifies the strength of the amplification channel, but the increase is mild.

Table 15: Sensitivity Analysis: Loan Balance Distribution and REO Discount

$\sigma = 1/52$	REO Disc. 10%	REO Disc. 20%	$\sigma = 1/65$	REO Disc. 10%	REO Disc. 20%
New York	(4.8%, 2.8%)	(6.2%, 3.4%)	New York	(6.6%, 4.0%)	(8.4%, 4.9%)
National	(22.3%, 11.9%)	(29.4%, 14.6%)	National	(30.0%, 17.3%)	(38.5%, 20.5%)
Las Vegas	(46.0%, 23.7%)	(60.2%, 28.7%)	Las Vegas	(62.0%, 34.6%)	(79.3%, 40.6%)
$\sigma = 1/78$	REO Disc. 10%	REO Disc. 20%	$\sigma = 1/91$	REO Disc. 10%	REO Disc. 20%
New York	(8.6%, 5.5%)	(10.8%, 6.5%)	New York	(10.9%, 7.2%)	(13.6%, 8.4%)
National	(39.3%, 23.9%)	(49.7%, 27.8%)	National	(51.0%, 32.2%)	(63.9%, 36.9%)
Las Vegas	(83.6%, 49.0%)	(105.6%, 56.4%)	Las Vegas	(115.4%, 69.5%)	(145.5%, 79.1%)

Note: This table shows comparative statics with respect to three important variables: σ, the Poisson probability of a renter becoming a buyer, the REO discount, and the loan balance distribution. The table shows how these variables affect the degree of amplification of the initial shock to prices created by adding foreclosures. The first entry in each pair is the percentage increase in the price index decline generated by defaults over and above that created by the initial shock. The second entry in each pair is the percentage increase in the total sales decline. For instance, when $\sigma = 1/52$ and the REO discount is 20%, the table has an entry of (60.2%,28.7%) for the 2006 Las Vegas loan balance distribution. Given a price index decline of 4.0% and volume decline of 4.4% from the initial shock alone, this indicates that the full price index and volume declines are respectively 6.4% and 5.7%.

Increasing the average length of time for which individuals who transition to renting stay out of the housing market leads to greater percentage decreases in both the price index and total volume. The effects of such a shock as small as the one we consider are relatively modest, but our key question remains the potential strength of the amplification channel when we allow for defaults. Table 15 reports the results from varying the REO discount and loan balance distribution in addition to σ. Rather than reporting levels, each entry reports the percentage amplification of the aggregate price index decline and sales decline generated by defaults over and above the decline created by the initial shock in a price-volume pair.

Table 15 shows that foreclosure spirals can significantly amplify an initial shock. The strength of this spiral grows as we increase the amount of time individuals who default spend out of the housing market due to a more persistent market tightness effect. The shape of the loan balance distribution also plays a critical role in determining the strength of the amplification. A greater proportion of individuals with high LTV ratios implies than a given initial shock will put a greater fraction of the market underwater. This leads to greater numbers of foreclosures, more powerful general equilibrium effects, and in turn even more foreclosures. This illustrates the fragility created by the combination of a housing bubble and reduced down payment requirements.

Table 15 also shows that a lower steady state REO discount dampens the spiral. This is in

part due to a compositional effect which ameliorates the effect of a given number of REOs on the price index. Additionally, though, the choosey buyer effect is weaker since waiting for an REO sale is no longer as attractive, and so the retail market does not freeze up as much. Finally, the amplification channel can generate significant *total* volume declines, greater than we saw in section 2.3 with exogenous defaults for a given shock size. The reason is that in the extended model, relative to section 2.3, the number of individuals who become locked-in during the downturn can be substantial. Note that this implies a greater REO share of vacancies which strengthens both the compositional effect and the choosey buyer effects and thus feeds back into further declines in both the overall and non-distressed price indexes.

2.7 Cross-MSA Quantitative Analysis

In order to assess quantitatively the role of foreclosures in the crisis and to test the model's performance, we calibrate the model to national and cross-MSA data described in section 2.2.

Because the size of the preceding bubble is the single best predictor of the size of the ensuing bust, we use a permanent shock to prices that operates through reducing flow utilities a to start the downturn. We assume income shocks last for 5 years, after which they gradually taper off. Recall that there are two parameters that are left to calibrate: σ, which controls how long the average renter takes to return to owner-occupancy, and the size of the permanent shock to prices. To calibrate these parameters, we use the aggregate national data and the cross-MSA data together. We first set a grid of σs. For each σ, we choose the nationwide permanent shock to prices so that with the nationwide loan balance distribution the model exactly matches the maximum log change in the national house price index. We then simulate the model for each MSA and for each value of σ by assuming that the permanent shock to prices in the MSA is equal to the nationwide permanent shock to prices multiplied by the relative size of the bubble in the MSA as measured by $\frac{\Delta \log\left(P_{03-06}^{MSA}\right)}{\Delta \log\left(P_{03-06}^{National}\right)}$. In other words, we assume that the relative amount of housing price appreciation from the bubble that is permanently lost is the same in each MSA.[80] We then calculate the unweighted sum of squared distances between the maximum change in the log price index in the data and the model for each value of σ and choose the σ with the minimum sum of

[80] For instance, the maximum $\Delta \log(P)$ in Las Vegas is was 1.52 times as big as the nation-wide price index. Below we find the permanent price decline for the nation is 21.5%. Thus the permanent decline in Las Vegas is 32.8%.

squared distances.

This methodology yields a sum of squared distances function that is a smooth and convex function of σ. Intuitively too high of a σ will cause the model to over-predict the size of price declines in high-LTV but low-bubble MSAs, while too low of a σ will cause the model to under-predict the size of price declines in high-bubble MSAs. The optimal σ is the one that does best across distribution of bubble size.

The sum of squared distances has a unique minimum, which corresponds to $\frac{1}{\sigma} = 1.05$ years out of the market for the average renter and $\Delta_{\max} P^{National} = 21.5\%$ (a log point decrease of .242).[81] A 21.5% price drop from the peak implies that about 2/3 of the price gains between 2003 and 2006 were permanently lost when the bubble burst.

Figure 21 shows the time series price, sales, the characteristics of distressed sales, probability of sale, foreclosures, and the mass of each type in the market for the resulting national simulation. The qualitative patterns closely match those described in the simpler model in Section 2.4. Recall that there is a one-to-one mapping between the unconditional probability of sale and time to sale, which move in opposite directions.

Figure 22 shows the results of the cross-MSA simulations by plotting the simulated results against the data in six panels that show the maximum change in log price, log retail price, log sales, log retail sales, REO share, and fraction foreclosed. In each figure, the 45-degree line is drawn in to represent a perfect match between the model and the data. The small dots represent MSAs while the large X represents the national calibration.

The calibration procedure does well in matching declines in the aggregate price index across the bubble-size spectrum, as indicated by the fact that the data points are clustered around the 45-degree line. As with the regressions, the extreme outliers are in greater Detroit (the two points in the lower right), which has had a large bust without a preceding boom, and Stockton (the point in the lower middle), which had a much larger bust than boom. Despite these outliers, we cannot reject a coefficient of one when regressing the simulated results on the data, and relative to a case with no default where the entire national price decline is permanent, adding default to the model increases the r-squared of the simulation by 20 percent.[82] These results suggest that default can

[81] 1.05 years out of the market for the average renter may be a bit on the short side, but because we have a constant Poisson probability of leaving renting, the distribution of times out of the market has a very thick tail.

[82] This r-squared is calculated as 1 minus the the squared distance between the simulated maximum log price decline

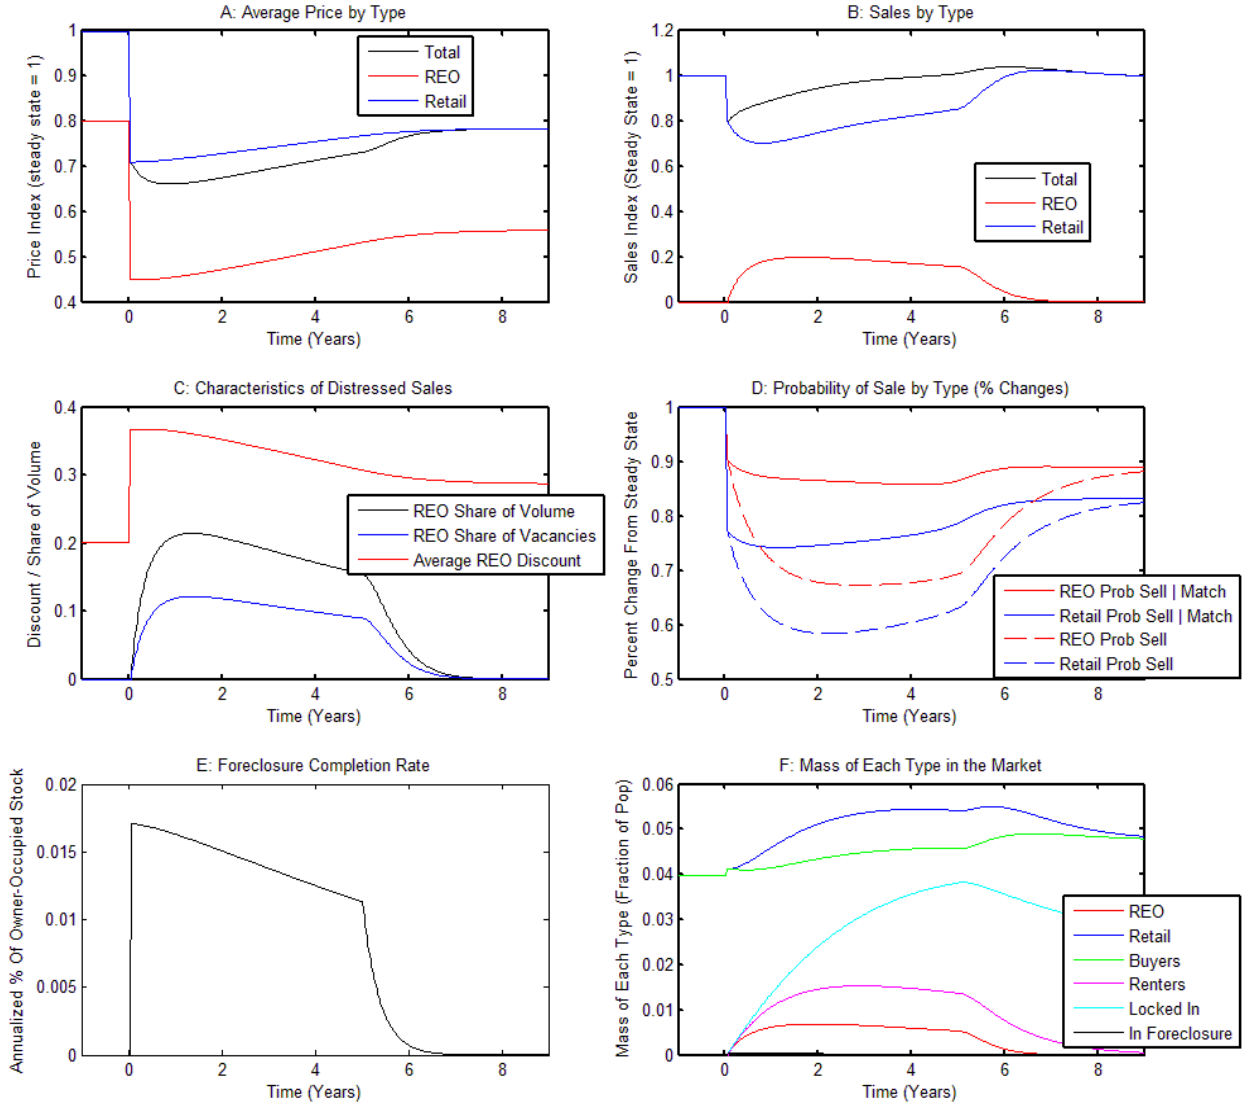

Figure 21: National Calibration With Permanent Price Drop

Note: This figure shows the results of the extended model calibrated to match the national and MSA data as in section 2.7. It uses a permanent shock to housing values and an LTV distribution corresponding to the national housing market. Panels A and B show the average price and sales by type, with pre-downturn price and volume normalized to 1. Panel C shows the REO discount, share of vacancies, and share of volume. Panel D shows the probability of sale conditional on a match and the unconditional probability of sale for each type with the pre-downturn probability normalized to 1. Panel E shows the annualized fraction of the owner occupied housing stock that is foreclosed upon at each point in time. Panel F shows the mass of each type of agent in the market. The overall results are similar to the qualitative results outlined previously.

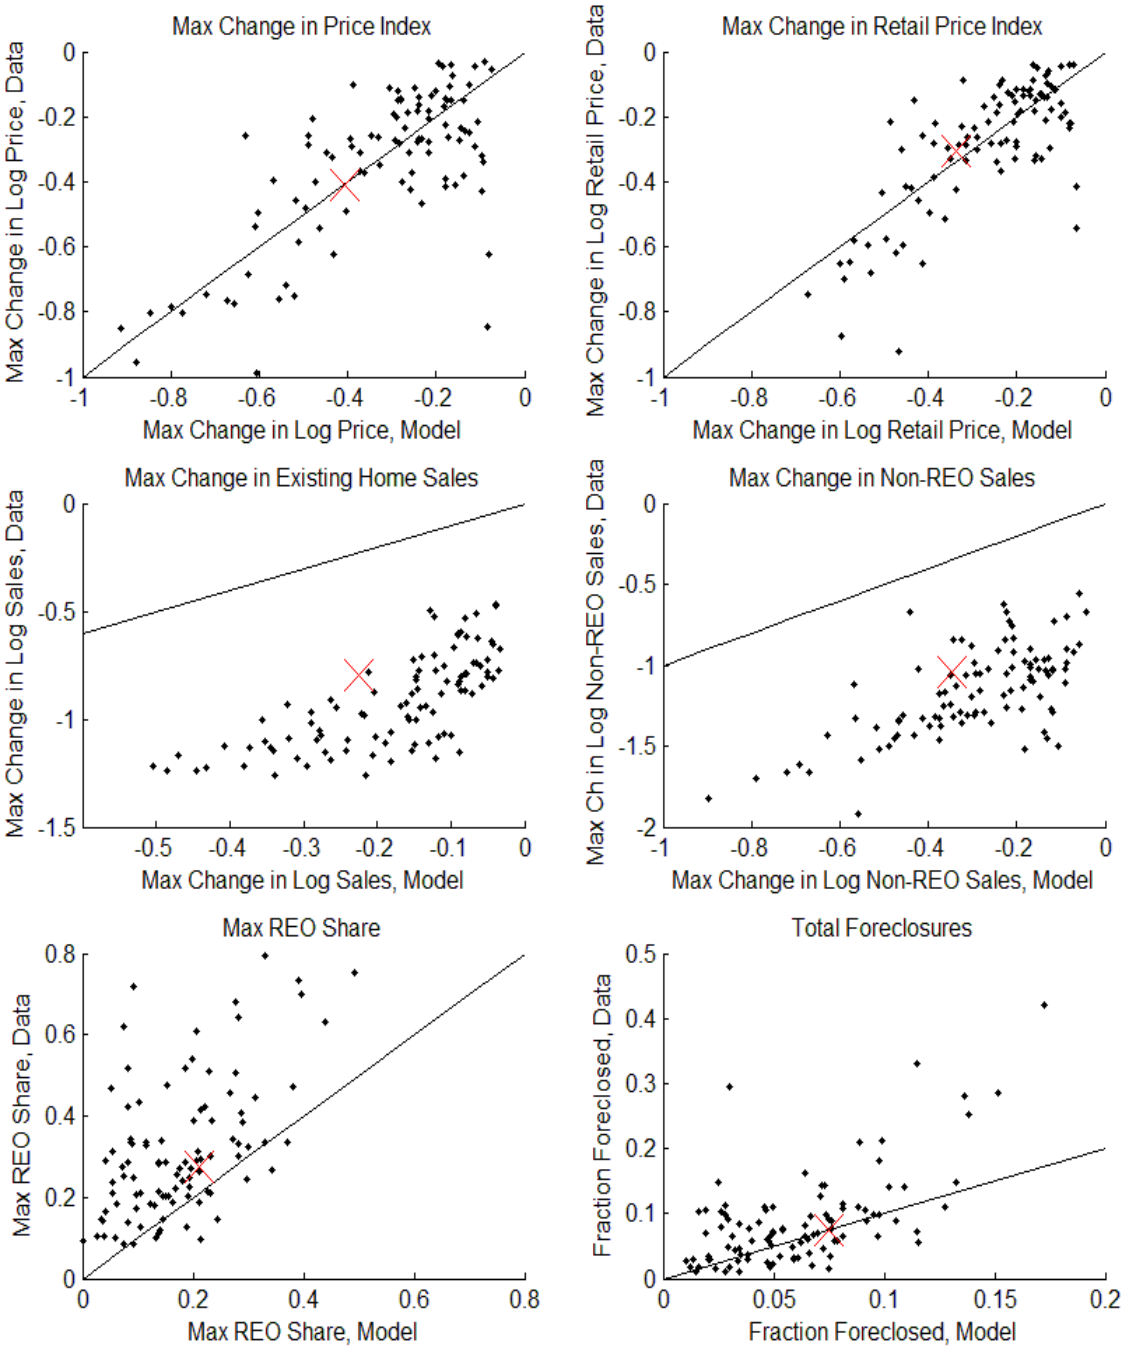

Figure 22: Cross-MSA Simulations vs. Data

Note: Scatter plots of data vs. simulation results for 97 MSAs in regression analysis. The red X represents the national simulation and each black dot is an MSA. The 45-degree line illustrates a perfect match between the model and the data. The variable being plotted shown in each plot's title. Data is fully described in appendix B.4. The calibration methodology described in text and appendix B.4. The figure shows the model performs well for prices and the number of foreclosures but is off by a constant for volume.

explain some of the nonlinearity in Figure 15.

More importantly, the model cannot be rejected for additional outcomes beyond the aggregate price index that was used for calibration. For the change in log retail price, the national data closely matches the model: for the most part, the points lie around the 45-degree line, with a few more exceptions where the model under-estimates the change in log retail price such as Stockton and Las Vegas. Nonetheless, a statistical test confirms that we cannot reject that the model matches the data. The model also comes close to matching the data for total foreclosures over 5 years and the model cannot be statistically rejected, although there are some extreme outliers where there were many more foreclosures than the model predicted. Again, some of these are in Greater Detroit and Stockton, but there are a few other hard-hit markets like Las Vegas and the Central Valley in California where the model under-predicts the number of foreclosures. Finally, although it is not in Figure 22, the national calibration predicts a maximum REO discount of 36.7%. This is slightly above the maximum foreclosure discount for Boston of 35.4% reported by Campbell et al. (2011), so the model can explain time variation in REO discounts.

However, as foreshadowed by the regressions, the model consistently under-predicts the decline in sales. In Figure 22, the data cluster roughly parallel to the 45-degree line for both retail and total sales, although this is only statistically significant for retail sales. This means that the model does a good job of capturing differences in the size of the maximum sales decline across locations but that volume has fallen nationwide for reasons beyond the model. Potential unmodeled forces reducing volume include the tightening of credit markets, credit constraints and losses on levered properties reducing the purchasing power of buyers (Stein, 1995; Ortalo-Magne and Rady, 2006), a decline in household formation and immigration, a reluctance on the part of retirees to sell their house in a down market, nominal loss aversion (Genesove and Mayer, 2001), increasing returns to scale in matching (Ngai and Tenreyro, 2013), and a reduction in the number of transactions by speculators who flip houses quickly. The cause of the massive decline in volume in the housing downturn is an important subject for future research.

Because it under-estimates the sales decline, the model also under-estimates the REO share in locations that had extremely high amounts of foreclosures, although when we include a constant in

and its counterpart in the data divided by the squared distance between the data and its mean. This is the r-squared of a regression of the model on the data without a constant.

a regression of the simulated results on the data we cannot reject a coefficient of one. Because the vast majority of sellers also become buyers, a decline in sales would strengthen the choosey buyer effect, as REO sellers would take up a greater fraction of the market. It would, however, have a much smaller impact on market tightness, because a reduction in the number of buyers and sellers would reduce both the numerator and the denominator. We expect the overall magnitude of the combined general equilibrium effect of foreclosures to be similar.

What do these figures imply about the quantitative extent to which foreclosures exacerbate housing downturns? In the national data, the permanent price decrease that would occur without default is 21.5% (.24 log points) and with default is 33.5% (.41 log points). This implies that the general equilibrium effects of foreclosures together with the compositional effects on the price index induced by a high REO share made the downturn 56% worse than it would have been in the absence of foreclosure. Equivalently, foreclosures account for 36% of the price decline. This figure is larger in MSAs with larger busts, more default, and a bigger price-default spiral.

The 56% figure, however, includes compositional effects and is thus not the best measure of how much the general equilibrium effects of foreclosure reduce the price a *retail* seller would get if they wanted to sell at the bottom of the market. This is the relevant price for determining negative equity and thus defaults. An alternate metric of the extent to which foreclosures exacerbate downturns, then, is the decline in the retail-only price index, which is 28.7% (.34 in log points) with default and 21.5% without default. The price decline in the retail market is thus 34% worse than it would have been in the absence of foreclosures.

Perhaps surprisingly, these quantitative results are not dramatically changed with an REO discount of 10% in steady state as suggested by Clauretie and Daneshvary (2009) and Zillow (2010). In this case, the same calibration procedure implies a permanent price decline of 22.4% and an average time out of the owner-occupied market for foreclosures of 1.3 years. Intuitively, with the compositional effect weakened by a smaller foreclosure discount, the calibration implies a slightly larger permanent price decline and a stronger market tightness effect. With a 10% steady state discount, the model implies that foreclosures exacerbate the aggregate price decline by 50%. See Appendix B.2.2 for details.

These magnitudes are larger than those implied by other papers. Mian et al.'s (2014) empirical study comes closest to our results. By comparing neighborhoods in states that require judicial

approval of foreclosure with neighborhoods just over a border in states that do not, they find that foreclosures were responsible for 20 to 30% of the decline in prices. Our analogous figure of 36% is only slightly higher, likely because we consider market-wide effects that comparing neighborhoods only partially picks up. Calomiris et al. (2008) use a panel VAR to analyze the effect of foreclosures on housing market equilibrium and find that foreclosures would reduce prices by 5.5 percentage points in a foreclosure wave, about half what we find. However, they simulate the impulse response to a wave of foreclosures without a bursting bubble that puts a substantial fraction of homeowners under water, which dramatically increases the size and length of the foreclosure wave. Using a calibrated macro model that focuses on how foreclosures can constrict credit supply, Chatterjee and Eyigungor (2011) find that foreclosures account for 16% of the overall price decline.[83]

2.8 Welfare and Policy Implications

2.8.1 Welfare

To evaluate welfare we adopt a utilitarian social welfare function that equally weights all agents. We can construct social welfare as the discounted sum of individual flow utilities:

$$W = \sum_{t=0}^{\infty} \beta^t \left\{ v_n(t) m_n + v_d(t) m_d + (l_1(t) + f(t)) (h_n^* + 1/\lambda) + \frac{f(t)}{\phi f(t) + 1} c + q(t) \right\}$$

where $q(t)$ follows the law of motion:

$$q(t+1) = (1-\gamma) q(t) + v_b(t) q_b(\mu(t)) \sum_m r_m(t) (1 - F(h_m(t))) (h_m(t) + 1/\lambda), \; q(0) = 0.$$

were $q(t)$ denotes the expected flow housing services generated to homeowners in l_0 at time t. We also assume that a foreclosure completion entails certain costs to the bank, such as legal fees and lost revenue from interest payments. A 2008 report by Standard & Poors estimates these costs of foreclosure in excess of the loss on the sale to be approximately $10,000, so that we set

[83] Our results also relate to an empirical literature that examines the effects of REO sales on the sale prices of extremely nearby houses. These papers typically find that a single REO listing reduce the prices of neighboring properties by 1%, with a nonlinear effect (Campbell et al, 2011). The literature is divided as to the mechanism. Anenberg and Kung (Forthcoming) arguing that REOs increase supply at an extremely-local level consistent with our market tightness effect, although at a very local level. By contrast, Gerardi et al. (2013(Gerardi, Willen, Rosenblatt, and Yao 2013)) argue that the owners of distressed property reduce investment in their home, effects for which we attempt to control. These papers typically include fine geographic fixed effects and consequently do not pick up the search-market-level effects that are substantial in our model.

$c = -10$. Finally, we suppose that individuals who receive a taste shock but are unable to move due to negative equity receive no flow utility from their mismatched house.

There are competing effects of foreclosures on social welfare. First, total welfare is decreased relative to the steady state since securing a foreclosure completion is costly, foreclosed homes are sold by REO sellers with higher holding costs, and all homes take longer to sell. More significantly, welfare is decreased by the fact that a number of homeowners receive no flow utility from housing because they are excluded from the housing market for a period of time due to default or locked into a house that does not suit their needs. Second, buyers who do participate in the market are on average purchasing homes which they value more than homes purchased in steady state. Because these buyers stay in these houses for a median of 9 years, this generates a substantial positive effect on welfare that is consistent with anecdotal evidence of the downturn being a "buyer's market." Note that the decline in prices which accompanies the downturn has no direct impact on welfare since it operates simply as a transfer from sellers to buyers that has no effect on social welfare. Ultimately, when all of the various effects net out, social welfare falls, but the decline is modest.

This likely understates policy makers' perception of the social impact of foreclosures. The welfare calculation uses a utilitarian framework and a high discount rate. Given that the downturn is temporary and a number of individuals actually benefit from the housing downturn in the form of increased housing services relative to steady state, it is not surprising we find a modest decline in welfare. However, it is still the case that a substantial mass of individuals are substantially worse off for several years. To the extent that policy makers adopt a Rawlsian short-term perspective, the social impact of foreclosures could be large.

Most importantly, by focusing only on the housing market, the model misses a number of other potential normative implications of foreclosures. As discussed by Iacoviello (2005), house price declines can have pecuniary externalities because a collapse in home prices destroys wealth in the form of home equity and can impede borrowing by households and firms, creating a financial accelerator effect similar to Kiyotaki and Moore (1997). Moreover, lock-in due to negative equity can impede labor mobility (Ferriera, Gyourko and Tracy, 2010) and exacerbate structural unemployment and can increase the effective risk faced by households since housing consumption is not adjustable (Chetty and Szeidl, 2007). Additionally, banks may be forced to realize substantial losses on foreclosed properties, which impacts their balance sheets through the well-documented

Table 16: Effects of Foreclosure Policies

Policy	Baseline	$\phi = 4000$	$\phi = 6000$	$7\% \rightarrow 4\%$	$L \downarrow \$2K$	$L \downarrow \$5K$
$\max \Delta \log (\mathrm{P})$	-.414	-.400	-.387	-.382	-.407	-.396
$\max \Delta \log (\mathrm{P}_{\mathrm{Retail}})$	-.343	-.344	-.347	-.327	-.339	-.334
$\max \Delta \log (\mathrm{Sales}_{\mathrm{Existing}})$	-.226	-.205	-.189	-.210	-.223	-.217
$\max \Delta \log (\mathrm{Sales}_{\mathrm{Retail}})$	-.352	-.322	-.295	-.314	-.339	-.321
$\max \frac{\mathrm{Sales}_{\mathrm{REO}}}{\mathrm{Sales}_{\mathrm{Existing}}}$	21.44	17.49	12.61	17.47	20.63	19.41
% Ever Foreclosed	7.55	7.54	7.48	6.10	7.27	6.86

Note: This table shows the impact of various policies on market equilibrium. The first and second rows show the maximum change in log aggregate and retail price, the third and fourth rows show the maximum change in log existing sales and retail sales, the fourth row shows the maximum REO share, and the sixth row shows the fraction of homes foreclosed upon. The first column shows the baseline, which is as in Section 2.7. The second and third columns show the effects of slowing down foreclosures, as $\frac{1}{\phi}$ foreclosures can be processed each week. $\phi = 4000$ and $\phi = 6000$ correspond to a maximum of 1.3 and .9 percent of the housing stock being foreclosed upon per year, respectively. The fourth column shows the effect of reducing interest rates for all homeowners from 7 to 4 percent, a generous estimate of the potential effects of refinancing mortgages. The sixth and seventh columns show the effects of reducing principal by \$2,000 and \$5,000 for all homeowners, corresponding to a \$100 billion principal reduction that is either untargeted or only targeted at under water homeowners.

net-worth channel in financial intermediation, potentially leading to problems in the interbank repo market, cash hoarding by banks, and a freezing up of credit. Finally, the presence of substantial numbers of foreclosed homes can have negative externalities on communities (Campbell et al., 2011) and can reduce residential investment, construction employment, and consumption (Mian et al, 2014).

While we leave detailed analyses of these important issues to future research, we believe the discussion in this section illustrates there is value in understanding the effectiveness of various policies in ameliorating the foreclosure crisis. We thus conduct a basic positive analysis of three policies that have been proposed that fit into our model: delaying foreclosure, refinancing mortgages at lower interest rates, and reducing principal. To assess the maximum potential impact of each policy, we introduce the policy at time 0. The results of our policy simulations, discussed in the following subsections, are shown in Table 16. We use the national loan balance distribution and compare the housing market under each policy to a baseline of no policy that is shown in column 1.

2.8.2 Delaying Foreclosure

A simple and low-cost policy that has been proposed is slowing down the pace of foreclosures. To incorporate sluggish foreclosure into our model, we assume that when a homeowner defaults the bank begins foreclosure proceedings but that only $\frac{1}{\phi}$ foreclosures can be processed by the system each week.[84] While in the foreclosure process, it is possible for prices to rise and the house to no longer be in negative equity. If this happens, the foreclosure "cures" and the homeowner lists their house as a normal seller but subsequently become a renter because of the liquidity shock they experienced. In table Table 16, we compare the baseline of $\phi = 0$ to cases when $\phi = 4{,}000$ and $\phi = 6{,}000$ so that the maximum annual pace of foreclosure is given by 1.3 percent per year and .9 percent per year, respectively. To further elucidate the effects of foreclosure backlogs, Figure 23 shows the aggregate and retail price indices and foreclosure starts and completions for the three values of ϕ.

A tighter foreclosure pipeline has different implications for the prices of retail homes versus the overall price index. In particular, the maximum decline in the overall price index falls because the compositional effects of foreclosure are weakened. However, the maximum retail price decline is greater and the price declines last for longer because the foreclosure crisis is extended: as panel B shows, even though foreclosure starts fall off after 5 years, with $\phi = 4,000$ the wave of foreclosures lasts over 6 years and with $\phi = 6,000$ it lasts nearly 9 years. This increases the duration of both the market tightness and choosey buyer effects which gets capitalized into lower retail prices.

The effect of foreclosure backlogs in our model is consistent with the argument that delaying foreclosures does not substantially prevent foreclosures in the long run and only draws out the pain. However, there may be benefits to delaying foreclosure that are not captured by the pure backlog story. For instance, if one expects household formation to pick up and boost demand in the near future, delaying foreclosures from a period of low demand to a period of higher demand could limit price declines. Similarly, slowing down foreclosures could cause banks to offer more mortgage modifications or short sales, reducing the number of delinquencies that result in a foreclosure.

In fact, the empirical evidence on states with judicial approval of foreclosure – in which backlogs

[84] Formally, we assume that if $f(t)$ homes are in the foreclosure process pending approval only $\frac{f(t)}{\phi f(t)+1}$ can be processed in a given period. We choose this function as a smooth approximation to $\min\left\{f, \frac{1}{\phi}\right\}$, which processes up to $\frac{1}{\phi}$ foreclosures each period. Such an approximation is necessary for the numerical implementation.

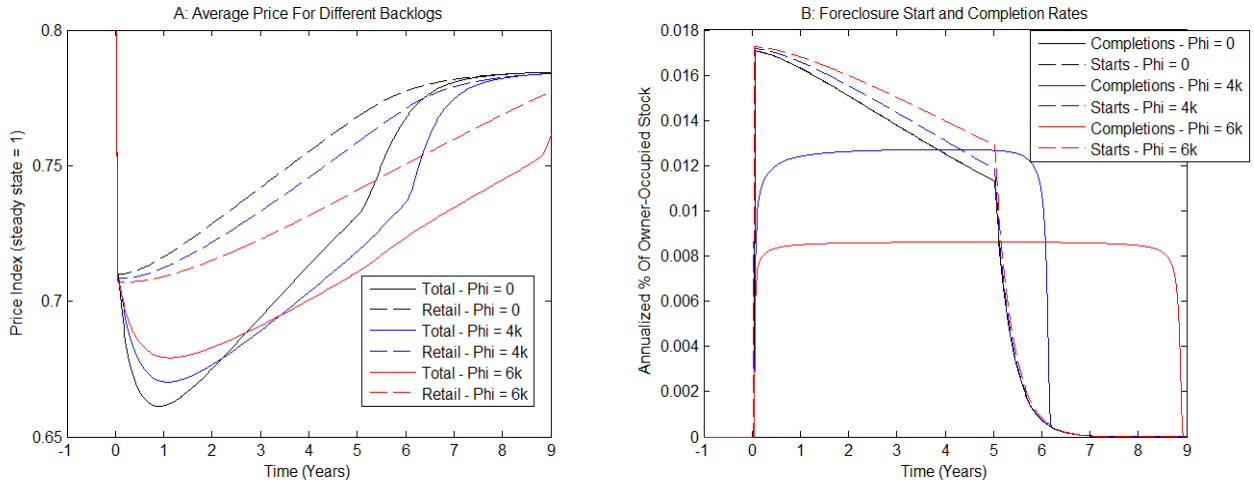

Figure 23: Policy: Various Sized Backlogs

Note: The figure shows the effect of prices and foreclosure start and completion rates for three different backlogs. The model is calibrated to the national calibration developed in section 2.7. The pre-downturn price level is normalized to one. One in ϕ houses can be foreclosed upon each week, so $\phi = 4{,}000$ corresponds to 1.3 percent of the housing stock being foreclosed upon per year and $\phi = 6{,}000$ corresponds to .9 percent. The figure shows that slowing down foreclosures extends the downturn and makes prices remain low for longer. Although the overall price index rises, this is because of a compositional effect and the retail price index falls.

are much larger (Mian et al., 2014) – suggests that slowing down foreclosures might reduce the incidence of foreclosure. Adding a judicial state dummy to regression (27) leads to a judicial dummy coefficient +.08 log points for the aggregate price index and +.05 for the non-distressed index even with a full set of controls, as shown in appendix B.3. Our model cannot generate such a dramatic price increase by adding a narrow foreclosure pipeline – the only way to get an effect of this order of magnitude is to reduce the incidence of foreclosures. The welfare effects of policies that limit the ability of lenders to foreclose by slowing down foreclosures are, however, unclear, as lenders may respond to a diminished ability to foreclose by increasing interest rates on mortgages or denying mortgages to credit-worthy borrowers.

2.8.3 Interest Rate or Payment Reductions

Another much-discussed policy that is being implemented with the Home Affordable Refinance Program is to refinance the mortgages of underwater borrowers, many of whom are stuck at extremely high interest rates due to an inability to refinance, at today's low interest rates. This could reduce

defaults because some individuals who are currently unable to meet their monthly payment may be able to pay a reduced monthly payment.

To simulate this intervention, we reduce γ_I, the hazard of default for individuals who are underwater, from 8.6 percent to 7.1 percent. Appendix B.4 uses an estimate of the effect of reducing monthly payments on defaults from Bajari et al. (2010) to show that this reduction in γ_I is equivalent to reducing interest rates from 7 to 4 percent – a generous estimate of what is possible purely through refinancing. The results are shown in column 4 of Table 16.

Although foreclosures still play an important role in exacerbating the downturn, refinancing has a substantial effect both because mechanically fewer foreclosures occur at a given level of negative equity and because the amplification mechanism is weaker. The size of these effects, however, depends critically on the effect of reducing interest rates on default. While we calibrate to the existing evidence from Bajari et al. (2010), their estimates are not causal. Understanding the impact of interest rate reductions on default is an important subject for future research.

2.8.4 Principal Reduction

The final policy we simulate is a $100 billion principal reduction.[85] The results are shown in Columns 5 and 6 of Table 16. Column 5 assumes that the government cannot target underwater homeowners gives every mortgage holder a $2,000 principal reduction. Column 6 assumes that the government can target the approximately 20 million individuals who have had negative equity during the crisis so that principal is reduced by $5,000 for each underwater homeowner.

Principal reduction provides a direct way to reduce negative equity and the price-default amplification. The targeted principal reduction has a significant ameliorating effect on the crisis, although it is not as effective as the interest rate reduction. The smaller principal reduction, however, has an effect that is much smaller. The government's ability to target homeowners in need is thus crucial to the effectiveness of principal reduction.

Beyond the government's ability to target under water homeowners, costs that we do not model may limit the effectiveness of principal reduction. Chief among these is moral hazard: if people expect that underwater mortgages will be bailed out with principal reductions, they may be more

[85] We assume that raising the funds for the intervention does not affect housing demand, housing prices, or the rate at which liquidity shocks occur.

likely to become delinquent on their mortgage. Similarly, strategic default may be elevated. The empirical relevance and size of such moral hazard effects is an important subject for future research.

2.8.5 Other Policies

Our policy simulations reveal the trade-offs faced by policy makers and the parameters that future research on anti-foreclosure policy should consider. In addition to the policies simulated here, there are a number of other policies that require a richer model to be given full justice and that we hope will be analyzed by future research.

First, a policy maker might try to stimulate additional buyer demand. To have a substantial effect on market tightness, one would have to stimulate entry by new homeowners. Such a policy is outside the scope of our model as it would require endogenous household formation or an endogenous buy-rent decision. Nonetheless, our model does suggest that any increase in new home ownership would have to be permanent; an intervention that boosts new home ownership for a few months at the expense of demand in subsequent months would not have a lasting effect. The short-lived effects of the 2009 new homeowner tax credit suggests that it is difficult to generate a long-lasting effect.

Second, our model cannot consider the conversion of owner-occupied housing to rental housing without an endogenous rent-buy margin. In particular, the conversion of REOs to rental properties has been discussed. While such a policy would reduce rental prices and REO inventories, with endogenous tenure choice it is possible that renting becomes much more appealing, drawing away buyers and further freezing up the owner-occupied market. There is also the potential for rent-seeking behavior by investors who seek to buy REO properties in bulk and convert them to rental homes.

2.9 Conclusion

This chapter argues that foreclosures play an important role in exacerbating housing downturns due to their general equilibrium effects. We add foreclosure to a simple search model of the housing market with two types of sellers by making two additional assumptions: banks selling foreclosed homes have higher holding costs than retail sellers and homeowners who are foreclosed upon cannot immediately purchase another home.

With these assumptions, foreclosures alter market behavior by reducing the number of buyers in markets, which makes sellers and particularly REO sellers desperate to sell, and by raising the probability that a buyer meets a REO seller who sells at a discount, which makes buyers more selective. Foreclosures also alter the composition of transactions, making the average sale look more like a foreclosure sale. These effects all create downward pressure on price but have opposing effects on volume as sellers want to sell faster but buyers are more choosey. Sales fall disproportionately in the retail market, helping to explain how foreclosures freeze up the market for non-distressed homes.

We then embed our basic model of the housing market in a richer model which allows for endogenous defaults and homeowner lock-in. We elucidate the potential for spirals in which foreclosures lower prices, putting more homeowners underwater, leading to more defaults and therefore even more foreclosures. A sensitivity analysis demonstrates that such a spiral can operate as a powerful amplification channel of shocks, especially when the proportion of homeowners in the market with high LTV ratios is high. A calibration of the full model to cross-market data is successful in matching both the average level of the price decline of the housing bust and a significant proportion of the cross-sectional variation in prices. The model matches the cross-sectional pattern of volume declines but is unable to fully account for the level. A quantitative exercise shows that foreclosures exacerbate the price declines in downturns on the order of 50 percent overall and 33 percent in the retail market.

An alternative explanation for the freezing up of the retail market during the housing bust is nominal loss aversion as documented by Genesove and Mayer (2001). The housing bubble may have created a reference point for homeowners such that when the bubble burst, they were not willing to sell for less than what they *perceived* the true value of their homes to be. If this were not the case for banks, the retail market would disproportionately freeze up. However, loss aversion would have to be extreme to explain a freezing up of the retail market for several years. Consequently, while it may not be able to fully account for the freezing up of the retail market, loss aversion may have played a role in the housing downturn and may be able to explain the volume declines that our model cannot capture. Note also that in a model in which nominal loss aversion is an operative channel, foreclosures may actually aid in price discovery.[86]

[86] Thanks to Ed Glaeser for this insight.

Credit constraints and capital losses on levered houses could also explain some of the freezing up of the retail market and the decline in volume that our model cannot explain. Ortalo-Magne and Rady (2006) present a model in which homeowners use equity extracted from their previous house to purchase their next house. With down payment requirements as in Stein (1995), moderate swings in housing prices can generate large swings the purchasing power of potential homeowners. This may cause some homeowners not to move at all, creating effective lock-in of non-underwater borrowers and helping to freeze up the retail market. The substantial decline in household formation is another factor that could explain the decline in volume that our model cannot explain.

Our analysis suggests several directions for future research. First, it would be interesting to endogenize the decision to enter the housing market. Chetty and Szeidl (2007) present an (S,s) model of consumption commitments based on the decision to move homes which could be embedded into a general equilibrium model of the housing market. This would allow an analysis of how market forces affect the decision to move and elucidate why sales remain depressed and more people are not taking advantage of the buyer's market by trading up. Second, an endogenous rent-buy margin would allow the rental market to be less segmented and would allow for analyses of several additional policies, as we describe in Section 2.8. Finally, the addition of supply considerations would allow an analysis of how the dynamics of new construction and conversion of owner-occupied housing to renter-occupied are affected by foreclosures.

Chapter 3:
Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities

3.1 Introduction

Macroeconomic models of fluctuations in hours of work over the business cycle or across countries imply much larger labor supply elasticities than microeconometric estimates of hours elasticities. Understanding this divergence is critical for questions ranging from the sources of business cycles to the impacts of tax policy on growth and inequality. Starting with the seminal work of Rogerson (1988) and Hansen (1985), one leading explanation of the divergence is the extensive margin response created by indivisible labor supply. If labor supply is indivisible, changes in tax or wage rates can generate large changes in aggregate hours by inducing extensive margin (participation) responses even if they have little effect on hours conditional on employment. In view of this argument, modern macro models are calibrated to match low micro estimates of intensive margin elasticities. However, the extensive margin elasticity is usually treated as a free parameter that can be calibrated purely to match macroeconomic moments.

We argue that the extensive margin elasticity should not be treated as a free parameter. Macro models should be calibrated to match micro estimates of extensive margin elasticities in the same way that they are calibrated to match micro estimates of intensive margin elasticities. The size of the extensive margin responses depends on the density of the distribution of reservation wages around the economy's equilibrium. The same marginal density that determines the impacts of macroeconomic variation on aggregate employment also determines the impacts of quasi-experiments such as tax policy changes on employment rates. Micro estimates of extensive margin elasticities can therefore be used to calibrate macro models.[87]

In this chapter, we assess whether existing calibrations of macro models are consistent with micro evidence on extensive margin responses. In doing so, we find that it is crucial to distinguish between two types of elasticities: Hicksian and Frisch. The Frisch (marginal utility constant) elasticity

[87] The distribution of reservation wages at the margin could vary across subgroups, potentially generating differences between micro and macro estimates of extensive-margin responses. As we explain below, observable heterogeneity in elasticities across subgroups reinforces our conclusions.

controls intertemporal substitution responses to temporary wage fluctuations and is therefore the relevant parameter for understanding labor supply fluctuations over the business cycle.[88] The Hicksian (wealth constant) elasticity controls steady-state responses to permanent wage changes and is therefore the relevant parameter for understanding differences in labor supply across countries with different tax systems. We use two approaches to comparing macro calibrations with micro evidence: simulations of quasi-experiments and a meta-analysis of micro elasticity estimates. Both approaches show that micro and macro evidence agree about Hicksian (steady state) elasticities but disagree about Frisch (intertemporal substitution) elasticities.

We begin by simulating the impacts of policy changes that generate exogenous changes in incentives to work in a standard macro model and comparing the predicted responses with the findings of microeconometric studies. We use Rogerson and Wallenius' (2009) [RW] calibrated model of life cycle labor supply as a benchmark model for this exercise. The RW model matches macro evidence by generating an intertemporal substitution elasticity of aggregate hours above 2 even when calibrated to generate a Frisch intensive-margin elasticity below 0.5. We simulate labor supply responses to three policies using this model: (1) a tax-free year in Iceland in 1987 studied by Bianchi et al. (2001), (2) a randomized experiment providing temporary subsidies for work to welfare recipients in Canada (Card and Hyslop 2005), and (3) the 1994 expansion of the Earned Income Tax Credit (EITC) for low-income individuals in the United States (Meyer and Rosenbaum 2001). The first two examples are ideal for identifying Frisch elasticities because they induce temporary variation in wage rates. Bianchi et al. (2001) find that employment rates in Iceland do indeed rise in 1987, but the increase is only one fifth as large as that predicted by the RW model. Similarly, the calibrated RW model predicts intertemporal substitution responses to the work subsidies in Canada that are nearly four times larger than what Card and Hyslop observe in their data. The third example – the EITC expansion – generates permanent variation in tax rates and thus is well-suited for identifying steady-state elasticities. The RW model performs better in matching the impacts of the EITC expansion on employment rates because it generates a Hicksian aggregate hours elasticity of approximately 0.7, resulting in steady-state impacts of taxes on labor

[88] The extensive margin Frisch elasticity is technically ill-defined because each agent is not at an interior optimum. We therefore define the Frisch extensive elasticity empirically as the impact of an infinitismal, temporary wage change on employment rates. This is the relevant elasticity for evaluating employment responses to business cycle fluctuations.

supply that are closer to micro estimates.

While our quantitative results rest on the particular assumptions of the RW model, our qualitative conclusions apply more generally. Any macro model that relies primarily on changes in labor supply to generate business cycle fluctuations must feature a large extensive margin Frisch elasticity. As a result, any such model will over-predict the response to temporary wage changes such as the tax holiday in Iceland and work subsidies in Canada. Intuitively, fluctuations in employment over the business cycle and the employment effects of quasi-experimental wage changes are both fundamentally determined by the same density of the reservation wage distribution at the margin irrespective of model specification. Thus, any labor supply model that fits the quasi-experimental evidence cannot generate large fluctuations in employment over the business cycle.

To explore whether the results of the three studies we consider in the simulations are representative of the broader empirical literature, we conduct a meta-analysis of quasi-experimental estimates of extensive margin elasticities. We summarize results from fifteen studies that span a broad range of countries, demographic groups, time periods, and sources of variation. These studies generally analyze changes in incentives for small subgroups of the population, permitting identification of labor supply elasticities that are not confounded by changes in equilibrium wage rates. Despite the great variation in methodologies, there is consensus about extensive margin elasticities. The mean extensive margin elasticity among the studies we consider is 0.28 and every estimate is below 0.43. The intertemporal substitution (Frisch) elasticity estimates for temporary policy changes turn out be quite similar to the steady-state elasticity estimates obtained from permanent policy changes. The small elasticities imply that most individuals are at a corner in their employment choices; that is, the density of individuals at the margin of employment is thin in practice.

We conclude our analysis by evaluating whether extensive margin elasticities of around 0.25 as suggested by micro evidence are adequate to reconcile the gap between micro and macro estimates of aggregate hours elasticities. To do so, we summarize micro and macro estimates of Hicksian and Frisch elasticities on both the extensive and intensive margins. We find that micro and macro studies agree about the steady-state impacts of taxes on labor supply. Both micro and macro studies imply Hicksian extensive margin elasticities around 0.2. And both micro and macro evidence are consistent with intensive margin elasticities around 0.3 once one accounts for frictions that may attenuate observed micro estimates (Chetty et al. 2011a, Chetty 2012). Prescott's (2004)

widely-cited cross-country dataset implies an aggregate hours (extensive plus intensive) Hicksian elasticity of 0.7, only slightly larger than micro estimates.[89] These findings indicate that labor supply responses to taxation could indeed explain much of the variation in hours of work across countries with different tax systems.[90]

On the intertemporal substitution margin, the limited existing evidence on intensive margin elasticities suggests that values around 0.5 are consistent with both micro and macro data. However, micro and macro estimates of extensive margin intertemporal substitution elasticities differ by an order of magnitude. Quasi-experimental estimates of extensive margin intertemporal substitution elasticities are around 0.25. In contrast, pure equilibrium macro models, in which employment fluctuations are driven purely by preferences, imply intertemporal substitution extensive margin elasticities in excess of 2. Hence, the puzzle to be resolved is why employment rates fluctuate so much over the business cycle relative to what one would predict based on the impacts of tax changes on employment rates – that is, why micro and macro estimates of the *Frisch* extensive margin elasticity are so different.[91] Even accounting for indivisible labor, micro studies do not support representative-agent macro models that generate Frisch elasticities above 1.

There are two potential concerns that one may have with using microeconomic estimates to calibrate macroeconomic models. The first is that heterogeneity in extensive margin responses complicates the mapping from micro estimates to macro elasticities that reflect economy-wide behavior.[92] This problem is compounded by the concern that micro studies sometimes exclude important subgroups that could matter for economy-wide extensive margin responses (Dyrda et al. 2012). In practice, however, heterogeneity across subgroups appears to strengthens our main conclusion about agreement on the Hicksian elasticity but disagreement on the Frisch elasticity. The

[89] Prescott reports an elasticity of approximately 3 in his paper. Importantly, this is a Frisch rather than Hicksian elasticity. Prescott implicitly maps the Hicksian elasticity of 0.7 implied by the data to a Frisch elasticity of 3 based on specific parametric assumptions. See Section 3.2 below for further details.

[90] Other factors, such as institutions or regulations, could also play a significant role in explaining cross-country hours differences (Alesina, Glaeser, and Sacerdote 2005). Our analysis does not rule out the importance of such factors. We simply show that micro estimates of labor supply elasticities are consistent with observed differences in aggregate hours across countries with different tax systems.

[91] Some progress has been made in recent years on this front: for instance, search and matching models with rigid wages (e.g. Hall 2009) can potentially match business cycle fluctuations with smaller extensive margin labor supply elasticities.

[92] Note that the same problem could in principle arise with intensive margin elasticities as well. Although macro models are often parametrized so that the intensive margin elasticity is constant by assumption, there is no economic reason for intensive margin elasticities to remain constant as wage rates change. Hence, if one is willing to use micro estimates to calibrate intensive margin elasticities, one should be equally willing to do so on the extensive margin as well.

heterogeneity in micro estimates of extensive-margin Hicksian elasticities mirrors the heterogeneity observed in macro studies of steady-state responses. For instance, both micro and macro studies indicate that extensive-margin elasticities are higher for subgroups that are less attached to the labor force, such as single mothers and individuals near retirement. However, heterogeneity magnifies the discrepancy between micro and macro estimates of intertemporal substitution elasticities. Most notably, employment rates fluctuate substantially over the business cycle even for prime-age males, which stands in sharp contrast with the near-zero micro extensive margin Frisch elasticity estimates for this group.

A second potential concern in mapping micro estimates to macro labor supply elasticities is that reduced-form micro studies may not directly identify the structural primitives of the reservation wage distribution that control extensive margin labor supply choices. This is particularly a concern if frictions prevent the labor market from clearing, as our analysis suggests. In a model with frictions, reduced-form micro elasticity estimates represent a convolution of the density of the reservation wage distribution at the margin and other structural parameters, such as the distribution of adjustment costs or search frictions or the degree of liquidity constraints. Importantly, the same reduced-form elasticities would also determine the impact of wage changes on labor supply over the business cycle in such an environment. Hence, micro estimates should continue to provide useful targets for calibrating macro models even though they do not identify the structure of preferences or other primitives necessary for normative analysis.[93] However, especially when reduced-form elasticities combine several structural parameters, they may not be stable across settings. Because of this instability, one should not seek to calibrate macro models to match any single estimate of a micro elasticity. Nevertheless, one can gauge the range of plausible magnitudes by pooling evidence from many different studies and settings as we do here. The fact that every quasi-experimental study we review finds elasticities significantly less than 0.5 casts doubt upon macro models calibrated with extensive margin elasticities above 1.

The chapter is organized as follows. Section 3.2 briefly reviews the existing literature on indivisible labor. In Section 3.3, we establish a terminology for the various elasticity concepts, as these

[93] Some micro studies attempt to strip out frictions by studying subgroups such as bike messengers or taxi drivers who can choose their daily labor supply more freely. However, it is not clear that these pure labor supply elasticity estimates are more relevant for macro calibrations. If the same frictions that constrain salaried workers from responding to tax changes also contrain their responses to fluctuations over the business cycle, then it is the observed reduced-form elasticity for the average worker that matters.

terms are often used in different ways in the existing literature. Section 3.4 reports simulations of the three quasi-experiments in the Rogerson and Wallenius (2009) model. Section 3.5 presents the meta-analysis of micro estimates. In Section 3.6, we compare micro and macro evidence on the intensive and extensive margins. Section 3.7 concludes. Details of the simulation methods and meta-analysis are given in the appendix.

3.2 Indivisible Labor: Background

Equilibrium macroeconomic models – in which differences in hours of work are driven by preferences – require large labor supply elasticities to explain the variation in hours of work over the business cycle and across countries with different tax regimes. In contrast, quasi-experimental microeconometric studies of the impacts of tax reforms on hours of work and earnings typically obtain elasticities close to zero for most groups except very high income earners.[94]

A large literature has posited that the discrepancy between micro and macro elasticities can be explained by indivisibilities in labor (e.g. Hansen 1985, Rogerson 1988, Cho and Rogerson 1988, Christiano and Eichenbaum 1992, Cho and Cooley 1994, King and Rebelo 1999, Chang and Kim 2006, Ljungqvist and Sargent 2006, Prescott, Rogerson, and Wallenius 2009, Rogerson and Wallenius 2009).[95] If individuals cannot freely choose hours of work or face fixed costs of entry, aggregate employment depends upon the distribution of reservation wages in the economy. If this distribution has substantial density at the margin – i.e., many individuals are indifferent between working and not working at prevailing wage rates – then a small reduction in wage rates could reduce aggregate hours of work significantly because many individuals will stop working. Yet the same change in wage rates may not affect hours of work conditional on employment very much, implying a small intensive margin labor supply elasticity. As a result, a model with large extensive margin elasticities and small intensive margin elasticities could match both the micro and macro evidence. Motivated by these results, modern macro models are calibrated to match micro estimates of intensive margin elasticities but typically calibrate the extensive margin elasticity

[94] Early estimates of intensive-margin elasticities include MaCurdy (1981), Altonji (1986), and Angrist (1991). Blundell and MaCurdy (1999) review this literature. Chetty (2012) and Saez, Slemrod, and Giertz (2012) summarize more recent quasi-experimental intensive margin elasticity estimates.

[95] The literature has taken two approaches to aggregation with indivisible labor supply: aggregation over states via employment lotteries (e.g. Hansen 1985, Rogerson 1988) or aggregation over time periods in a lifecycle model (e.g. Mulligan 2001, Ljungqvist and Sargent 2006, Prescott, Rogerson, and Wallenius 2009). The micro evidence on extensive margin responses we review here is most easily interpreted through the modern life cycle models.

purely to match macroeconomic moments (King and Rebelo 1999, Rogerson and Wallenius 2009, Ljungqvist and Sargent 2011).

In parallel with the development of macro models of indivisible labor supply, a large microeconometric literature has recognized the importance of the extensive margin in the analysis of labor supply. Ashenfelter (1984) and Heckman (1984) discuss the importance of extensive margin labor supply choices in the analysis of aggregate fluctuations. Heckman and Killingsworth (1986) and Heckman (1993) review the literature on labor supply models that explicitly model participation decisions. More recent research has estimated extensive margin elasticities using quasi-experimental methods.

However, macro models have not been calibrated to match micro evidence on extensive margin elasticities. One complication in performing such a calibration is that extensive margin elasticities vary with the wage rate unless the density of the reservation wage distribution happens to be uniform. Hence, any micro estimate of an extensive margin elasticity is necessarily local to the wage variation used for identification. However, this argument does not justify treating the extensive margin elasticity as a free parameter for two reasons. First, if the micro estimates are identified using variation similar to that used in macroeconomic comparisons, one will obtain the appropriate local elasticity relevant for macro calibrations. Second, the same problem arises when calibrating macro models with micro estimates of intensive margin elasticities, insofar as elasticities will only be constant on the intensive margin if utility happens to produce a constant-elasticity labor supply function. We revisit this issue in Section 3.6 and show that, if anything, observable heterogeneity in elasticities reinforces the conclusions drawn below.

3.3 Terminology

It is helpful to establish some conventions about terminology given the various elasticity concepts discussed in this chapter. We distinguish between elasticities based on the margin of response (extensive vs. intensive) and the timing of response (intertemporal substitution vs. steady state). There are four elasticities of interest: steady-state extensive, steady-state intensive, intertemporal extensive, and intertemporal intensive. Each of these four elasticities can be estimated using both micro (quasi-experimental) and macroeconomic variation. We use the terms "micro" and "macro" elasticities exclusively to refer to the source of variation used to estimate the elasticity.

The elasticity of aggregate hours – the relevant parameter for calibrating a representative agent model – is the sum of the extensive and intensive margin elasticities, weighted by hours of work if individuals have heterogeneous preferences (Blundell, Bozio, and Laroque 2013).

The macro literature uses the term "macro elasticity" to refer to the Frisch elasticity of aggregate hours and "micro elasticity" to refer to the intensive-margin elasticity of hours conditional on employment (e.g. Prescott 2004, Rogerson and Wallenius 2009). We use different terminology here for two reasons. First, the intensive-margin is no more "micro" than the extensive margin; both are determined by household-level choices and both have been estimated using micro data. Second, and more importantly, the Frisch elasticity is critical for understanding business cycle fluctuations in models where aggregate hours fluctuations are purely driven by labor supply, but it is not the relevant parameter for evaluating the steady-state impacts of differences in taxes across countries. The Frisch (marginal utility constant) elasticity controls intertemporal substitution responses to temporary wage fluctuations, while the Hicksian (wealth constant) elasticity controls steady-state responses and the welfare consequences of taxation (MaCurdy 1981, Auerbach 1985).[96]

The distinction between Hicksian and Frisch elasticities is quite important in practice. Prescott (2004) reports that cross-country differences in aggregate hours imply an elasticity of 3 in a representative-agent model, whereas Davis and Henrekson (2005) estimate an elasticity of 0.33 using similar data. The difference arises primarily because Prescott reports a Frisch elasticity whereas Davis and Henrekson report a Hicksian elasticity. Regressing log hours on log tax rates in Prescott's data yields a Hicksian elasticity of 0.7, as shown in 25a below. Prescott maps this estimate of the Hicksian elasticity into a value for a Frisch elasticity based on parametric assumptions about utility and the wealth-earnings ratio. When utility is time-separable, the Frisch (ε^F) and Hicksian (ε^H) elasticities are related by the following identity (Ziliak and Kniesner 1999, Browning 2005):

$$\varepsilon^F = \varepsilon^H + \rho(\frac{d[wl]}{dA})^2 \frac{A}{wl},$$

where ρ is the elasticity of intertemporal substitution (EIS), $\frac{d[wl]}{dA}$ is the marginal propensity to earn out of unearned income, and $\frac{A}{wl}$ is the ratio of assets to earnings. The reason that Prescott obtains

[96] The Hicksian elasticity determines the impact of taxes in steady-state if government revenues are returned to the consumer as a lump sum, as commonly assumed in representative-agent macro models. If revenues are not returned to consumers, tax changes have income effects and the Marshallian elasticity becomes the relevant parameter.

a much larger value of ε^F than ε^H is that the parametric utility specification he uses produces large values of $\frac{A}{wl}$ and $\frac{d[wl]}{dA}$. However, microeconometric evidence shows that income effects on labor supply are much smaller than those produced by the Prescott utility specification (Holtz-Eakin, Joulfaian, and Rosen 1993, Imbens et al. 2001). Under a utility specification that matches empirical estimates of the mean values of $\frac{d[wl]}{dA}$ and $\frac{A}{wl}$, the Frisch elasticity is only slightly larger than the Hicksian elasticity because the difference between the two elasticities is proportional to the income effect squared $(\frac{d[wl]}{dA})^2$ (Chetty 2012, Table III).[97]

3.4 Simulations of Quasi-Experiments in the RW Model

We evaluate whether macro models with indivisible labor are consistent with micro evidence on extensive margin responses by focusing on the Rogerson and Wallenius (2009) model. The RW model is a leading example of recent models of indivisible labor that aggregate over individuals by time-averaging over the life cycle, as in Ljungqvist and Sargent (2006). The model is well-suited for our purposes because it features both an extensive and intensive margin of labor supply. RW calibrate their model to show that small intensive-margin micro elasticities are consistent with a large Frisch elasticity of aggregate hours. We adopt the parameters chosen by RW and simulate the impacts of policy changes analyzed in three prominent microeconometric studies.[98]

Setup. RW analyze an overlapping-generations model in which a unit mass of agents is born at each instant and lives for one unit of time. An individual who supplies $h(a) \in [0,1]$ hours at age a produces $e(a) \times \max\left\{h(a) - \bar{h}, 0\right\}$ efficiency units of labor, where $e(a) = 1 - 2(1-e_1)\left|\frac{1}{2} - a\right|$ is a tent-shaped life-cycle productivity profile and $\bar{h} > 0$. Complete asset markets lead to perfect consumption smoothing. With log utility over consumption, each generation solves

$$\max_{c,h(a)} \log(c) - \alpha \int_0^1 \frac{h(a)^{1+\gamma}}{1+\gamma} da \text{ s.t. } c = (1-\tau) \int_0^1 e(a) \max(h(a) - \bar{h}, 0) da + T$$

where τ is the tax rate and T is a lump-sum tax rebate that balances the government's budget.

[97] Subsequent studies calibrate models to match Prescott's Frisch elasticity of 3, but choose a different functional form for utility and wealth-earnings ratios (e.g. Trabandt and Uhlig 2011). The conclusions drawn by these studies – e.g. that reductions in tax rates would increase tax revenue – might differ had they directly matched the steady state elasticity of 0.7 implied by Prescott's data.

[98] On the intertemporal substitution margin, we sought to maximize the model's chance of fitting the data by analyzing the two studies that obtain the largest intertemporal elasticity estimates among those considered in our meta analysis (Table 17). On the steady-state response, we chose a representative study of a well-known policy (the Earned Income Tax Credit) to show that the model is consistent with typical micro estimates.

The model can be solved analytically as described in RW and in Appendix C.4. Because wages are paid per efficiency unit, individuals have low hourly wage rates at the beginning and end of their lives and find it optimal not to work at those points. This generates an extensive margin of participation over the life cycle. The convex disutility over hours of work generates an intensive margin hours response to changes in wage rates as well. RW normalize the price of output to 1 and assume a constant-returns-to-scale production technology, so changes in tax rates have no impact on pre-tax wages and prices. Accordingly, the quasi-experiments we simulate also hold pre-tax wages and prices constant, as the studies on which they are based typically analyze the impacts of differential changes in incentives for relatively small subgroups of the population.

RW calibrate the parameters α, e_1, and $\bar{h}$ to match empirically observed values for the fraction of life worked (f), the maximum hours worked per week over the life cycle ($h_{\max}$), and the wage rate at retirement relative to the maximum wage rate over the life cycle ($w_R/w_{\max}$). Following RW, we set $h_{\max} = 45\%$ (45 hours per week) and $w_R/w_{\max} = 1/2$. We set f to match the aggregate employment rate in the period prior to each policy experiment we consider. The parameter γ controls the Frisch elasticity of labor supply, as in standard life cycle models (Card 1990). We set $\gamma = 2$ to obtain an intensive margin Frisch elasticity of $\varepsilon_{\text{INT}} = \frac{1}{\gamma} = 0.5$, consistent with the microeconometric evidence summarized below; we show in Appendix C.1 that setting $\varepsilon_{\text{INT}} = 0.25$ yields similar results.[99] For each of the three tax policy changes simulated below, we choose the model's remaining parameters $\{\alpha, e_1, \bar{h}\}$ to match the moments $\{h_{\max}, w_R/w_{\max}, f\}$ under the tax system prior to the tax change.[100] In all three cases, the calibrated model generates an intertemporal substitution elasticity for aggregate hours above 2 despite having an intensive margin intertemporal substitution elasticity of only 0.5, consistent with RW's main result.[101] As in RW, we assume that each agent lives for 60 years (corresponding to average adult working lives) and simulate each quasi-experiment by changing the tax rate for the number of periods in the model

[99] RW show that the intertemporal elasticity of aggregate hours in their model is not sensitive to the intensive-margin intertemporal elasticity. They therefore calibrate α, e_1, and $\overline{h}$ to match the three moments conditional on various values of γ.

[100] In one of the simulations, the welfare demonstration in Canada, a small enough fraction of the population is employed prior to the intervention that fitting $w_R/w_{\max} = 1/2$ would require negative productivity at certain points in the life cycle. Consequently, for that simulation, we set $e_1 = 0$, generating $\frac{w_R}{w_{\max}} = .615$.

[101] We calculate this and all other Frisch elasticities by simulating the impact of a small, temporary tax change in the RW model. This direct calculation of the Frisch elasticity differs from the values reported by RW. RW report aggregate hours Frisch elasticities for a stand-in household whose behavior matches the aggregate steady-state properties of their economy. However, this stand-in household's behavior does not necessarily match the aggregate intertemporal substitution properties of the RW model.

that correspond to the duration of the tax policy change in the data.[102]

To simulate the impacts of unanticipated tax changes, we must specify how the lump sum rebate T changes for each agent. To simplify aggregation, we assume that each generation receives a lump-sum rebate equal to the taxes they pay at each instant in time.[103] We ignore heterogeneity in the tax system across individuals and set τ equal to the average tax rate for the subgroup analyzed (which is relevant for extensive margin decisions).

Experiment 1: Tax Holiday in Iceland. In 1987, Iceland suspended its income tax for one year as it transitioned from a system under which taxes were paid on the previous year's income to a system where taxes were paid on current earnings. In 1987, individuals paid tax on income earned in 1986; in 1988, individuals were taxed on income earned in 1988, and thus income in 1987 was untaxed. The average tax rate was 14.5% in 1986, 0 in 1987, and 8.0% in 1988 (Bianchi et al. 2001). Although this tax change could also produce a change in labor demand due to a general equilibrium impact on wage rates, the tax holiday had no impact on labor supply for individuals with low initial tax rates (Bianchi et al. 2001, Figure 9). This implies that the general equilibrium feedback on wage rates was negligible, so the aggregate employment response can be interpreted as a labor supply elasticity.[104] We simulate the tax reform in Iceland in the RW model under the assumption that the tax system remains stable prior to 1986 and after 1988. The reform was announced in late 1986, so we model the tax change as an unanticipated change at the start of 1987. The average employment rate in the three year period prior to the reform is $f = 79.2\%$, which implies that individuals work for 47.5 years in the model. The single-year tax reduction thus comes close to the ideal experiment for identifying a Frisch elasticity of reducing tax rates for an infinitesimal fraction of the working life.

Figure 24a plots annual changes in employment rates (the employment rate in year t minus the employment rate in year $t-1$) around the reform, demarcated by the vertical line. The Icelandic administrative records analyzed by Bianchi et al. (blue squares) show a modest but significant increase in employment rates in 1987 followed by a sharp dip in 1988, consistent with intertemporal

[102] To characterize high frequency dynamics precisely, we simulate the model with at least 100 periods per year in all cases; see Appendix C.4 for details.

[103] Tax policy changes affect each generation differently because they are at different points in the lifecycle when the change occurs.

[104] Stated differently, the differential response for workers who experienced larger changes in tax rates can be interpreted as a pure labor supply elasticity that nets out changes in wage rates. Bianchi et al.'s analysis reveals that this differential impact is similar to the aggregate impact we simulate here.

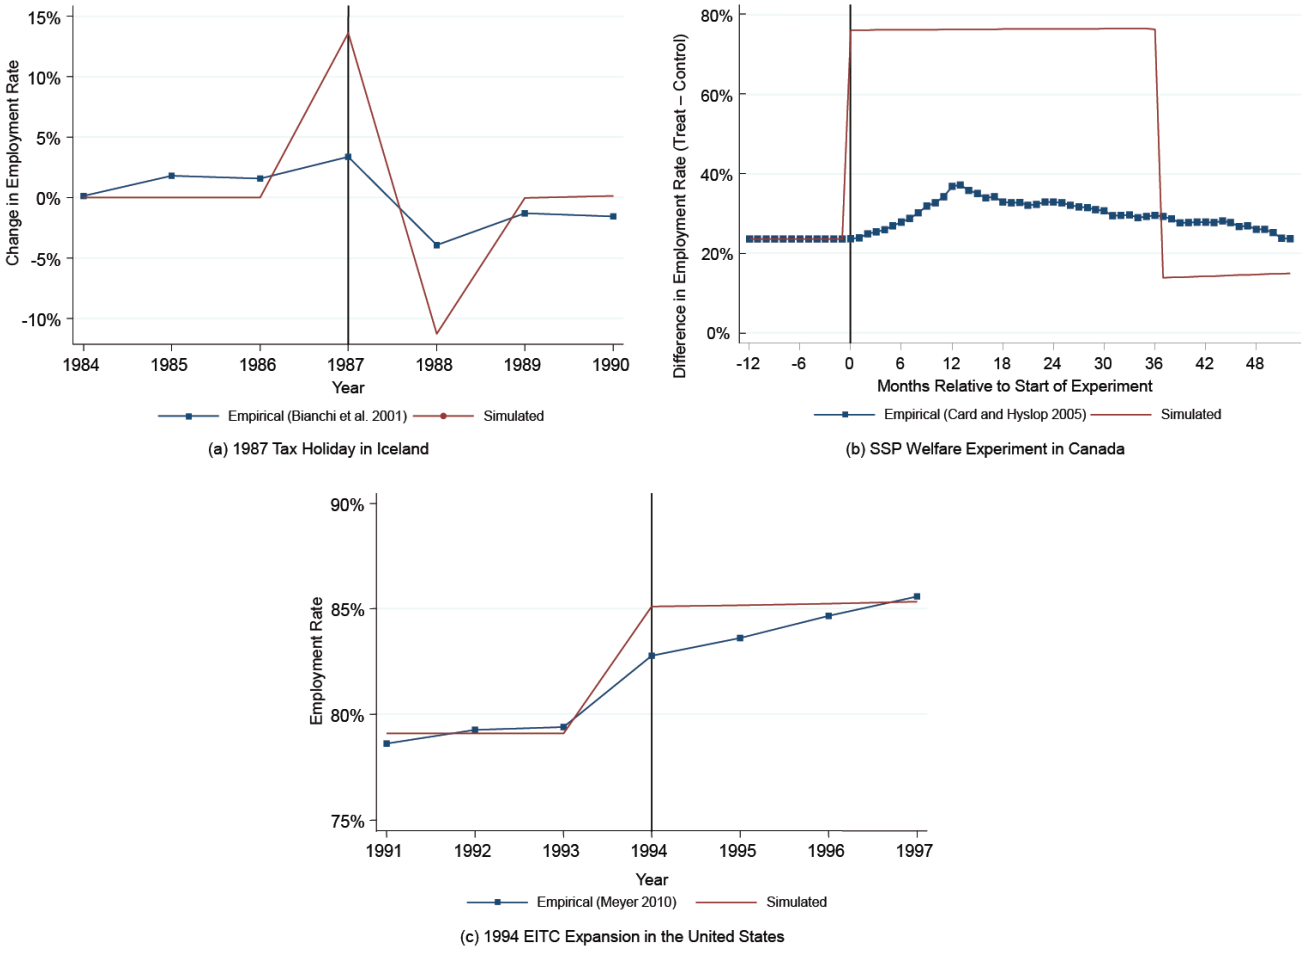

Figure 24: Impacts of Tax Changes on Employment Rates: Simulations vs. Data

Notes: Each panel shows the impact of an unanticipated change in incentives to work on employment rates. The red dashed series shows the impact predicted by the calibrated Rogerson and Wallenius (2009) model, while the blue solid series shows the impact observed in the data. Panel (a): Iceland suspended its income tax for one year in 1987. Average tax rates in Iceland changed from 14.5% in 1986 to 0% in 1987 and then 8.0% in 1988. Following Bianchi et al. (2001), we define the employment rate as the fraction of weeks worked in a given year in the adult population. This panel plots annual changes in employment rates. Panel (b): The Canadian SSP demonstration randomly assigned a group of welfare recipients a wage subsidy for 36 months in the early 1990s. Individuals in the control group faced an effective average tax rate of 74.3% for working full time at the minimum wage, while individuals in the treatment group faced an effective average tax rate of 16.7%. Following Card and Hyslop (2005), we plot the difference in monthly employment rates between the treatment and control groups. We add the observed control group mean at the start of the experiment (23.5%) to the difference for scaling purposes. Simulated employment rates are the fraction of individuals aged 16 to 46 working in a given month, reflecting the age distribution of the SSP treatment group (see Appendix A). Panel (c): The EITC expansion in the US in 1994-5 lowered average tax rates net of taxes and transfers for single mothers from 50.8% in 1992 to 43.6% in 1996. Meyer (2010, Figure 2) reports annual employment rates for single women using CPS data. We plot the employment rates of single mothers adjusted for observables and time trends as in Meyer (2010); simulated employment rates are reported for individuals aged 16 to 46.

substitution. The impact predicted by the RW model (red circles) is an order of magnitude larger than the observed impact. In the data, employment is 3 percentage points higher in 1987 relative to 1988, but the RW model predicts that it would be 13.5 percentage points higher. The model generates a much larger spike in employment because the fraction of cohorts that are close to being indifferent between working and staying out of the labor force is large. The temporary increase in the wage rates therefore induces a large group of agents to work. Note that it is precisely this mechanism – having a large fraction of individual near the margin – that allows the RW model to generate a large Frisch elasticity for aggregate hours and thus explain fluctuations in aggregate hours over the business cycle.

Experiment 2: SSP Welfare Demonstration in Canada. The Iceland analysis focuses on employment changes in the aggregate economy, which are relevant for understanding business cycle fluctuations but may mask substantial heterogeneity across groups. Ljungqvist and Sargent (2006), Rogerson and Wallenius (2007), and others emphasize that certain groups of the population – such an individuals near retirement or those with low wage rates – are likely to exhibit particularly large extensive margin responses and drive the change in aggregate hours. To evaluate whether the model's predictions are more accurate for these more elastic subgroups, we consider a policy experiment targeted at welfare recipients who frequently transition in and out of the labor force.

In the early 1990s, the Canadian government conducted the Self Sufficiency Project (SSP) to test whether a temporary earnings subsidy could induce welfare recipients to start working. The project was a randomized experiment involving over 5,000 single parents who had been on welfare for at least one year. Half the individuals (the treatment group) were given a large subsidy if they worked more than 30 hours per week. The subsidy lasted for 36 months.[105] Under the prevailing welfare system in Canada, welfare payments were reduced dollar-for-dollar with earnings above a low baseline level. As a result, a single parent with one child in the control group faced an effective average tax rate of 74.3% when moving from no work to full-time work (see Appendix C.1). In contrast, an individual in the treatment group faced an effective average tax rate of 16.7% for the same change. The employment rate during the month the experiment began was $f = 23.5\%$.

[105] Individuals were given up to one year to start working and the 36 month period began after they started to work. This feature of the program generated an incentive to establish eligibility for the subsidy by working within the first year, accentuating the intertemporal substitution incentive. We ignore this feature of the program in our simulation by assuming that the subsidy starts immediately after random assignment. This simplification biases the size of the employment increase predicted by our simulation downward.

Card and Hyslop (2005) use survey data to calculate employment rates at a monthly frequency for 53 months starting from the month of random assignment. Figure 24b plots monthly employment rates after the experimental intervention began. The series in blue squares shows the difference in employment rates for the treatment group relative to the control group (Card and Hyslop, Figure 3a), with the model the SSP experiment as a tax reform that lowers the tax rate from $\tau = 74.3\%$ to $\tau = 16.7\%$ for a three year period, after which the tax rate reverts to $\tau = 74.3\%$. The pre-experiment employment rate of 23.5% is added to the difference to facilitate interpretation of the scale. The data show that the subsidy had a substantial impact: employment rates rise by approximately 14 percentage points in the treatment group relative to the control group a year after the subsidy was introduced. These employment gains fade away after the subsidy expires, consistent with intertemporal substitution.

The series in red circles in Figure 24b shows the corresponding impacts predicted by the RW model. Because the sample analyzed by Card and Hyslop consists primarily of younger individuals (less than 2.5% of the sample is over age 50), we report simulated employment rates for individuals in the first half of the life cycle (ages 16-46). The impacts predicted by the calibrated model – an employment increase of 52.8 percentage points one year after the subsidy is introduced – are again substantially larger than what is observed in the data. Hence, even for subgroups that are closer to the margin of entering or exiting the labor force and are therefore more elastic, the RW model significantly over-predicts extensive margin responses.

One may be concerned that liquidity constraints attenuate the degree of intertemporal substitution in the low-income population treated by the SSP. The estimated elasticity therefore may not directly identify preference parameters in the RW model. However, as noted above, the same liquidity constraints should also affect employment responses to business cycle fluctuations in wage rates. Hence, the reduced-form response estimated by Card and Hyslop is still informative about the magnitude of labor supply fluctuations over the cycle for this subgroup.

Experiment 3: Earned Income Tax Credit in the U.S. The last policy change we consider – the expansion of the EITC in 1994 analyzed by Meyer and Rosenbaum (2000, 2001) and Meyer (2010) – is a permanent tax change whose impact is determined by the Hicksian rather than the Frisch elasticity.[106] The EITC expansion lowered average tax rates (including implicit taxes generated

[106] If the tax change is not rebated to the consumer as a lump sum, its impact depends on the uncompensated

by the phase-out of transfers) from 50.8% in 1992 to 43.6% in 1996 for single mothers (Meyer and Rosenbaum 2000, Table 2).[107] Roughly half of the expansion occurred in 1994. For simplicity, we model the tax change under the assumption that the change occurs immediately at the start of 1994, ignoring the phase-in of the reform. We also assume as above that the tax system remains stable prior to 1994. The average employment rate for the single mothers is $f = 79.1\%$ in the three years preceding the reform.

Figure 24c shows the employment rates of single mothers around the 1994 reform using data from Meyer (2010, Figure 2). The series in blue circles shows employment rates for single mothers with 1 or 2 children, adjusted for time trends and changes in observables as in Meyer (2010) (see Appendix C.1 for details). The labor force participation rate of single mothers rose from 79.6% in 1993 to 85.8% in 1997 after the EITC expansion was fully phased in. The RW model predicts a 6.0 percentage point increase in employment rates on impact and an additional 0.3 percentage point rise over the subsequent 5 years. The impact predicted by the model is thus very similar to the observed impact.

The RW model performs much better in predicting the impacts of the EITC expansion than the preceding experiments because it predicts much smaller steady-state responses than intertemporal substitution responses. Intuitively, a permanent change generates a much lower elasticity because all generations increase their labor supply at the point in their life cycle when they are most productive, smoothing the aggregate response across time. With a temporary change, every generation has an incentive to work when net-of-tax wage rates are high, resulting in a large Frisch elasticity.[108] In the RW model, a large mass of cohorts is at the margin with respect to a temporary tax change or wage fluctuation because individuals do not have strong preferences over when they work during their lives. However, in any given period, a much smaller fraction of individuals within each cohort are at the margin with respect to a permanent change in incentives.

Together, the simulations highlight two results that we develop further below. First, the

(Marshallian) elasticity rather than the Hicksian elasticity. In practice, microeconometric estimates of income effects are quite small (Holtz-Eakin, Joulfaian, and Rosen 1993, Imbens, Rubin, and Sacerdote 2001), suggesting that the impact of the EITC change is well approximated by the Hicksian elasticity.

[107] The changes in average tax rates in Meyer and Rosenbaum (2000) take into account conurrent changes in benefits from welfare and Medicaid. See Appendix A for details.

[108] Although the SSP welfare demonstration was temporary, a three-year subsidy actually covers a sizable fraction of the working life. The responses to the experiment are therefore determined by a combination of Hicksian and Frisch elasticities. Together, the Iceland and SSP simulations demonstrate that the RW model over-predicts responses both at very short and medium-term frequencies.

extensive margin elasticities required to explain the sharp fluctuations in aggregate hours over the business cycle are far larger than micro estimates. Second, micro and macro evidence are in much closer alignment about the steady-state impacts of taxes on labor supply.

Although the quantitative results of our simulations depend to some extent upon the parametric choices made by RW, we expect these lessons to apply more broadly. Generating a large macro Frisch elasticity by having a large fraction of individuals who are nearly indifferent between working and not working is precisely what delivers predictions about how temporary tax changes affect employment rates that contradict the data. A macro model calibrated to match micro estimates of extensive margin intertemporal substitution elasticities would no longer generate large Frisch elasticities for aggregate hours.

3.5 Meta-Analysis

In this section, we evaluate whether the three quasi-experiments considered above are representative of the broader literature by conducting a meta-analysis of extensive margin elasticity estimates. Although several papers have reviewed intensive margin elasticities (e.g. Pencavel 1986, Blundell and MaCurdy 1999, Chetty 2012), we are not aware of a meta-analysis of quasi-experimental estimates of extensive margin elasticities.

We focus on reduced-form studies that use changes in tax policies or long-term wage trends for identification rather than structural studies that exploit variation in wage rates at the individual level to fully identify a structural model. Keane and Rogerson (2010) argue that obtaining consistent structural estimates from wage variation over the life cycle requires accounting for a broad range of factors such as human capital accumulation (Imai and Keane 2004), credit constraints (Domeij and Floden 2006), and uninsurable risks (Low 2005). Moreover, structural models typically rely on strong exclusion restrictions for identification.[109] The quasi-experimental studies we consider here exploit variation that is orthogonal to wage rates and thus are more robust to the biases emphasized by Keane and Rogerson. The exclusion restriction underlying these studies is

[109] Common instruments for wage rates include nonlinear age and time trends (Kimmel and Kniesner 1998) or interactions of education and experience (Gourio and Noual 2009) conditional on individual fixed effects. Keane (2010) uses years of schooling as an instrument for the wage to identify an elasticity in Eckstein and Wolpin's (1989) classic structural model. The exclusion restrictions for these instruments are that employment rates do not vary with age conditional on wage rates or that individuals with different levels of education do not have different employment trajectories over their lifecycle. If factors that predict high wage rates also predict high latent tastes for work, the elasticity estimates would be biased upward.

Table 17: Extensive Margin Elasticity Estimates From Quasi-Experimental Studies

Study	Elasticity	Standard Error	Population and Variation
A. Steady State (Hicksian) Elasticities			
1. Juhn, Murphy, and Topel (1991)	0.13	0.02	Men, skill-specific trends, 1971-1990
2. Eissa and Liebman (1996)	0.30	0.10	Single Mothers, U.S. 1984-1990
3. Graversen (1998)	0.24	0.04	Women, Denmark 1986 tax reform
4. Meyer and Rosenbaum (2001)	0.43	0.05	Single Women, U.S. Welfare Reforms 1985-1997
5. Devereux (2004)	0.17	0.17	Married Women, U.S. wage trends 1980-1990
6. Eissa and Hoynes (2004)	0.15	0.07	Low-Income Married Men & Women, U.S. EITC expansions 1984-1996
7. Liebman and Saez (2006)	0.15	0.30	Women Married to High Income Men, U.S. tax reforms 1991-97
8. Meghir and Phillips (2010)	0.40	0.08	Low-Education Men, U.K. wage trends, 1994-2004
9. Blundell, Bozio, and Laroque (2013)	0.30	n/a	Prime-age Men and Women, U.K., tax reforms 1978-2007
Unweighted Mean	**0.25**		
B. Intertemporal Substitution (Frisch) Elasticities			
10. Carrington (1996)	0.43	0.08	Full Population of Alaska, Trans-Alaska Pipline, 1968-83
11. Gruber and Wise (1999)	0.23	0.07	Men, Age 59, variation in social security replacement rates
12. Bianchi, Gudmunndsson, and Zoega (2001)	0.42	0.07	Iceland 1987 zero tax year
13. Card and Hyslop (2005)	0.38	0.03	Single Mothers, Canadian Self Sufficiency Project
14. Brown (2009)	0.18	0.01	Teachers Near Retirement, California Pension System Cutoffs
15. Manoli and Weber (2011)	0.25	0.01	Workers Aged 55-70, Austria severance pay discontinuities
Unweighted Mean	**0.32**		

Notes: This table reports elasticities of employment rates with respect to wages, defined as the log change in employment rates divided by the log change in net-of-tax wages. Where possible, we report elasticities from the authors' preferred specification. When estimates are available for multiple populations or for multiple specifications without a stated preference among them, we report an unweighted mean of the relevant elasticities. See Appendix C.2 for details on sources of estimates.

that the differential changes in tax rates across groups is not correlated with unobserved determinants of employment rates, typically a weaker assumption than those required forfull identification of a structural model.[110]

Table 17 summarizes extensive margin elasticity estimates from fifteen quasi-experimental studies. The calculations underlying the estimates and standard errors are described in Appendix C.2. We calculate the extensive margin labor supply elasticity as the change in log employment rates divided by the change in log net-of-tax wage rates. Employment rates are typically defined as working at any point during the year, though there are some differences across studies as described in the appendix. We use the authors' preferred estimate whenever possible. For studies that do not report such an estimate, we construct elasticities from reported estimates of changes in participation and calculations of the change in net-of-average-tax wage rates. We use the delta method to calculate standard errors in such cases.

The studies summarized in Table 17 report labor supply elasticities for various countries and subgroups using many different sources of variation. Yet the elasticity estimates exhibit substantial consensus. The elasticity estimates range from 0.13 to 0.43, with an overall unweighted mean across the fifteen studies of 0.28. To obtain further insight into the key patterns, we divide the studies into two groups – steady-state and intertemporal substitution – based on the type of variation they use for identification.

The first panel in Table 17 shows steady-state (Hicksian) elasticities identified from permanent wage changes resulting from tax reforms or long term trends in wage rates across regions or skill-groups.[111] The simplest empirical designs (e.g. Eissa and Liebman 1996) use difference-in-differences approaches, while more recent studies (e.g. Meghir and Philips 2010) combine multiple reforms over time that affect individuals differently. The mean elasticity across the nine studies that estimate steady-state elasticities is 0.25.

The second panel in Table 17 summarizes results from studies that exploit temporary wage changes to identify intertemporal substitution (Frisch) elasticities. Some of these studies exploit

[110] Keane (2010(Keane 2010)) and Keane and Rogerson (2010) review structural estimates and find larger values than the quasi-experimental estimates summarized below. It would be useful to simulate the impacts of tax policy changes in these structural models to understand why their predictions differ from the reduced-form evidence.

[111] Some of the studies in Panel A of Table 17 do not fully account for income effects and thus obtain estimates that are closer to Marshallian elasticities than Hicksian elasticities. However, we can still conclude from the mean estimates in Panels A and B of Table 17 that the Hicksian elasticity is between 0.25 and 0.32 because the Hicksian is bounded by the Marshallian and Frisch elasticities (MaCurdy 1981).

temporary tax changes such as the Iceland tax holiday discussed above or temporary increases in labor demand, such as Carrington's (1996) analysis of the effect of the Trans-Alaska Pipeline on Alaska's labor market. Other studies analyze the impact of anticipated variation in wages generated by pension schemes on retirement behavior. For instance, Gruber and Wise (1999) correlate employment rates of adults near retirement with the implicit tax generated by social security systems across OECD countries. Their analysis implies an elasticity of 0.23. Brown (2009) and Manoli and Weber (2011) estimate elasticities using the bunching of retirements around the kinks in the budget set created by discontinuities in pension systems. The small elasticities found by these studies implies that the fraction of individuals who are "at a corner with respect to the decision to retire" (Ljungqvist and Sargent 2011) is quite large in practice.

The mean estimate of the intertemporal substitution elasticity across the six studies in Panel B is 0.32, only slightly larger than the estimates of steady-state elasticities in Panel A. The similarity between Hicksian and Frisch elasticities is consistent with evidence that income effects are not large enough to produce a substantial difference between intertemporal substitution and steady-state responses.[112]

The elasticity estimates vary across subgroups in correspondence with their mean employment rates, as is well known from prior work (Heckman 1993, Keane and Rogerson 2010). Groups that have the weakest attachment to the labor force, such as single mothers or older workers near retirement, are the most elastic on the extensive margin (e.g. Meyer and Rosenbaum 2001, Gruber and Wise 1999). Among prime-age males, high rates of labor force participation and low aggregate hours elasticities (which combine the intensive and extensive margins) have led researchers to conclude that the extensive margin response is likely to be quite small (see e.g., Hausman 1985 and Juhn, Murphy, and Topel 1991). This is why most of the studies in Table 17 focus on groups with relatively low participation rates. Hence, the mean extensive margin elasticity in the population as a whole is likely to be below the unweighted mean across the studies in Table 17 of 0.28.

The heterogeneity in elasticities across subgroups implies that there is no single value of the extensive margin elasticity that can be used across applications. For instance, a recession or tax policy change that affects prime-age males may generate smaller employment responses in the

[112] This does not imply that income effects are small in magnitude. Because the gap between the Frisch and the Hicksian is proportional to the *square* of the income effect, even sizable income effects $\frac{d[wl]}{dA}$ produce a small gap between the Frisch and Hicksian elasticities; see Chetty (2012(Chetty 2012)) for details.

Table 18: Extensive Margin Elasticity Estimates From Quasi-Experimental Studies

		Intensive Margin	Extensive Margin	Aggregate Hours
Steady State (Hicksian)	micro	0.33	0.25	0.58
	macro	0.33	0.17	0.50
Intertemporal Substitution	micro	0.54	0.32	0.86
(Frisch)	macro	[0.54]	[2.77]	3.31

Notes: Each cell shows a point estimate of the relevant elasticity based on meta analyses of existing micro and macro evidence. Micro estimates are identified from quasi-experimental studies; macro estimates are identified from cross-country variation in tax rates (steady state elasticities) and business cycle fluctuations (intertemporal substitution elasticities). The aggregate hours elasticity is defined as the sum of the extensive and intensive elasticities. Macro studies report intertemporal aggregate hours elasticities but do not always decompose these values into extensive and intensive elasticities. Therefore, the estimates in brackets show the values implied by the macro aggregate hours elasticity if the intensive Frisch elasticity is chosen to match the micro estimate of 0.54. See Appendix C for sources of these estimates.

macroeconomy than a change in incentives that affects other groups. The estimates in Table 17 should therefore be interpreted as a rough guide to plausible targets for calibration: they suggest that extensive margin elasticities around 0.25 are reasonable, while values above 1 are not.

3.6 Comparing Micro and Macro Estimates

The micro evidence points to Frisch and Hicksian extensive margin elasticities around 0.25. Does this estimate generate aggregate hours elasticities consistent with macro evidence? The answer to this question depends on the size of intensive margin elasticities because aggregate hours elasticities combine extensive and intensive elasticities. We therefore begin by summarizing the micro and macro evidence on both extensive and intensive margins in Table 18. The sources and calculations underlying these estimates are described in Appendix C.3. The rows of Table 18 consider steady-state (Hicksian) vs. intertemporal substitution (Frisch) elasticities, while the columns compare intensive margin (hours conditional on employment) and extensive margin (participation) elasticities. Within each of the four cells, we report micro and macro estimates of the elasticity based on (unweighted) means of existing studies. We also calculate aggregate hours elasticities – the parameter relevant for calibrating representative agent models – by summing the extensive and intensive elasticities.[113]

[113]For micro studies, this calculation requires that preferences are homogenous across the population. If groups that work few hours have higher extensive elasticities, as suggested by existing evidence, this calculation yields an upper bound on the aggregate hours elasticity (Blundell, Bozio, and Laroque 2011).

It is important to note that there are wide confidence intervals associated with each of the point estimates in Table 18, as well as ongoing methodological disputes about the validity of some of the underlying studies (see e.g., Saez, Slemrod, and Giertz 2012). Therefore, the estimates should be treated as rough values used to gauge orders of magnitude: differences of 0.1 between elasticity estimates could well be due to noise or choice of specification, while differences of 1 likely reflect fundamental discrepancies. We consider the evidence on steady-state and intertemporal elasticities in turn.

Steady-State. On the extensive margin, our rough estimate of the steady state elasticity from the micro literature is the mean of the estimates in Panel A of Table 17, which is 0.25. On the intensive margin, Chetty (2012) presents a meta-analysis of micro estimates of Hicksian elasticities and reports a mean value of 0.15 (Chetty 2012, Table 1). However, Chetty argues that these elasticities are attenuated by optimization frictions: the small tax changes used to identify micro elasticities do not generate substantial changes in hours because the adjustment costs agents have to pay to change hours outweigh the second-order benefits of reoptimization. Chetty develops a bounding method of recovering the underlying structural elasticity relevant for evaluating the steady-state impacts of taxes. Pooling the 15 studies he analyzes (Table 1, Panels A and B), he obtains a preferred estimate of the structural intensive margin Hicksian elasticity of 0.33.[114]

Macro steady-state estimates are obtained from comparisons across countries with different tax regimes. Nickell (2003) and Davis and Henrekson (2005) find extensive steady-state elasticities of 0.13 and 0.14, respectively, by regressing log employment-population ratios on log mean net-of-tax rates across countries. Prescott's (2004) tax data coupled with measures of labor force participation rates implies an extensive steady-state elasticity of 0.25 (see Appendix C.3). Our rough estimate of the steady state extensive margin elasticity from the macro literature is the mean of the estimates from these three studies, which is 0.17. Davis and Henrekson (2005) estimate a steady-state intensive elasticity of 0.20 by regressing log hours conditional on employment on log net-of-tax rates. As noted above, Prescott's (2004) data produces a steady-state aggregate hours elasticity

[114] Our proposed elasticities of 0.33 on the intensive margin and 0.26 on the extensive margin may appear to contradict the common view that tax changes have smaller short-run effects on the intensive margin than extensive margin. Chetty (2012) argues that the structural intensive margin elasticity relevant for long-run comparisons is larger than the structural extensive margin elasticity once one accounts for frictions. In particular, he shows that frictions attenuate observed extensive margin elasticities much less than intensive margin elasticities because the utility gains from reoptimizing are first-order on the extensive margin and second-order on the intensive margin.

of 0.7; subtracting the extensive margin macro elasticity of 0.25 produced from Prescott's data therefore implies an intensive steady-state elasticity of 0.46. The mean intensive margin elasticity implied by Prescott and Davis and Henrekson's analysis is 0.33, which we use as our estimate of the macro intensive margin elasticity.

We conclude that micro and macro estimates of steady state aggregate hours elasticities match once one accounts for extensive margin responses and optimization frictions.[115] Figure 25a illustrates the agreement by plotting log of hours per adult vs. log net-of-tax rates using the same cross-country data as Prescott (2004). The solid green line shows the best fit to Prescott's data, which generates a Hicksian elasticity of 0.7 as noted in Section 3.2. The dashed red line shows the relationship predicted by our preferred estimate of the micro aggregate hours elasticity of 0.58 from Table 18 (with the intercept chosen to match the mean values in the data). The similarity of the two lines illustrates the concordance between micro and macro estimates of steady-state elasticities.

Intertemporal Substitution. On the extensive margin, our preferred micro estimate of the intertemporal elasticity is the mean of the estimates in Panel B of Table 17, which is 0.32. On the intensive margin, there is less quasi-experimental evidence on intertemporal substitution elasticities. Bianchi et al. (2001) find an intensive-margin elasticity from the Iceland reform of 0.37 (see Chetty (2012) for the elasticity calculation using Bianchi et al.'s estimates). Pistaferri (2003) reports a Frisch intensive elasticity of 0.7 using microdata on expectations about wages. The mean of these two estimates is 0.54. It is not surprising that these estimates of the intensive Frisch elasticity are only slightly larger than our preferred estimate of the intensive Hicksian elasticity of 0.33. Chetty (2012) shows that the Frisch elasticity must be less than 0.47 given a Hicksian elasticity of 0.33 in an intensive-margin model with balanced growth and an intertemporal elasticity of substitution of consumption below 1. Utility specifications that generate a Frisch elasticity that is much larger than the Hicksian elasticity are inconsistent with micro estimates of income effects and elasticities of intertemporal substitution of consumption.

Equilibrium macro models identify intertemporal substitution labor supply elasticities from fluctuations in hours over the business cycle. Most macro studies calibrate representative agent models and therefore report only intertemporal elasticities of aggregate hours. The intertemporal

[115] The similarity between micro and macro estimates may be surprising given the institutional and regulatory differences across countries (Alesina, Glaeser, and Sacerdote 2005). However, institutions and regulations can partly be interpreted as sources of optimization frictions, which we account for using Chetty's (2009) bounding procedure.

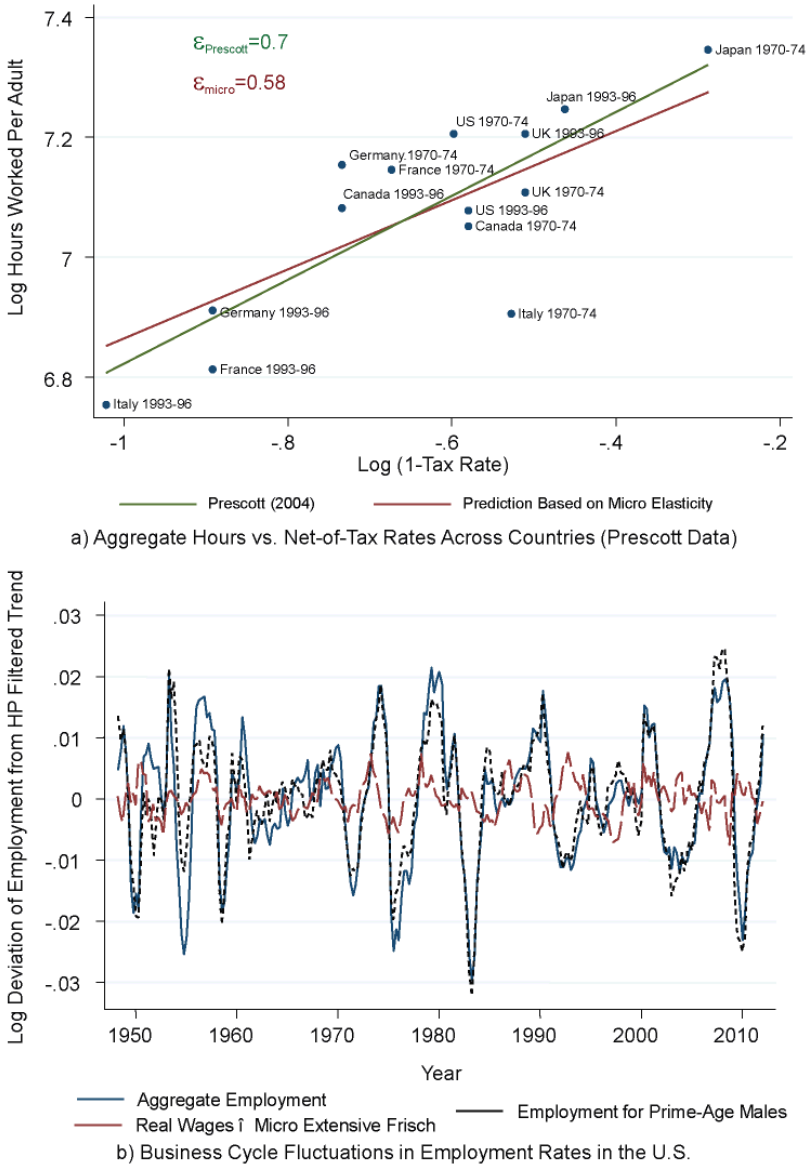

Figure 25: Micro Predictions Versus Macro Data

Notes: Panel A plots log hours worked per adult vs. log of 1 – average tax rate using data from Prescott (2004) across countries and time periods described in Appendix C. The data imply an aggregate hours Hicksian elasticity of .7, as shown by the solid green best fit line. The dashed red line is drawn through the mean of the x and y values with a slope of 0.58, in accordance with the aggregate hours micro elasticity from Table 18. Panel B plots the log deviation of employment from a Hodrick-Prescott filtered trend for the United States from 1948 to 2008. The data is taken from the Bureau of Labor Statistics and available at http://www.bls.gov. The solid blue line is generated using seasonally adjusted quarterly data on employment tabulated from the Current Population Survey, series LNS12000000Q. The raw data was Hodrick-Prescott filtered using a smoothing parameter of 1600. The dotted black line is taken from the same source for men ages 25-54, series LNS12000061Q. The dashed red line is a projected employment series based on fluctuations in real wages. Real wages are measured as real hourly compensation for the nonfarm business sector, Bureau of Labor Statistics series PRS85006153. To generate the projection, real wages are Hodrick-Prescott filtered using a smoothing parameter of 1600 and multiplied by the micro extensive margin Frisch elasticity of 0.32 from Table 17.

aggregate hours elasticity required to match business cycle data ranges from 2.6 to 4 in real business cycle models (Cho and Cooley 1994, Table 1; King and Rebelo 1999, p975). Table 18 reports the mean intertemporal aggregate hours elasticities implied by these numbers, 3.31.[116] Micro estimates imply a Frisch elasticity of aggregate hours of 0.86, well below the values implied by RBC models.

The few available decompositions of macro aggregate hours elasticities into extensive and intensive margins suggest that macro estimates are roughly in alignment with micro estimates on the intensive margin. Business cycle fluctuations in hours conditional on employment account for only 1/6 of the fluctuations in aggregate hours at an annual level (Heckman 1984). Given that elasticities of 2.6 to 4 fit the fluctuations in aggregate hours, we infer that intensive Frisch elasticities around 0.43 to 0.66 would match macro evidence in RBC models. These values are modestly larger than the intensive intertemporal elasticity of 0.5 implied by micro evidence.

In contrast, macro evidence sharply contradicts micro estimates of the extensive intertemporal elasticity. The fact that employment fluctuations account for 5/6 of the fluctuation in aggregate hours suggests that extensive elasticities of 2.18 to 3.33 would be needed to match the data in standard RBC models.[117] If the RBC models considered in Table 18 were calibrated to match an intensive intertemporal elasticity of 0.54, they would require extensive intertemporal elasticities of 3.31-0.54 = 2.77 on average to match aggregate hours fluctuations. This value is an order of magnitude larger than all of the micro estimates in Table 17.

We conclude that extensive labor supply responses are not large enough to explain observed fluctuations in employment rates at business cycle frequencies. This result is illustrated in Figure 25b. The solid blue line in the figure shows fluctuations in employment rates over the business cycle in the U.S. It plots the log deviation of employment (measured using household surveys) from a Hodrick-Prescott filtered trend. The dashed red line shows predicted employment fluctuations due to labor supply using our preferred micro estimate of the extensive margin Frisch elasticity of 0.32. The prediction is constructed by multiplying the Frisch elasticity of 0.32 by log deviations

[116] An earlier version of this table (Chetty et al. 2011b) included an estimate 1.92 from Smets and Wouters (2007) when computing the macro estimate of the intertemporal substitution elasticity. While Smets and Wouters report an estimate of 1.92, in personal correspondence they noted that the correct elasticity implied by their model is the reciprocal of the reported estimate, $\frac{1}{1.92} = 0.52$. This elasticity is much lower than traditional models because of a large number of frictions including wage and price rigidities, which make the Smets and Wouters paper significantly different from the pure equilibrium macro models discussed here. We thank Susan Yang for pointing out this correction.

[117] Cho and Cooley (1994) decompose the aggregate hours elasticity in their RBC model into intensive and extensive margins using a different methodology. Their analysis generates an extensive Frisch elasticity of 1.61.

in real wages from a Hodrick-Prescott filtered trend. The fluctuations in the data are much larger than the prediction based on micro evidence, illustrating that fluctuations in labor supply account for only a small share of observed employment fluctuations over the business cycle.

The size of the fluctuations in the micro prediction may be attenuated because of composition bias in the BLS wage series. Barsky, Solon and Parker (1994) argue that actual wages are approximately twice as volatile as observed wages because of changes in the composition of employed workers. With this adjustment, one would need an aggregate hours elasticity of $3.31/2 = 1.66$ to fit the macro data. While accounting for composition bias helps reduce the gap substantially, it does not fully reconcile the discrepancy between the macro business cycle data and predictions based on micro evidence.

Heterogeneity. As emphasized by Dyrda et al. (2012), macro models may not exactly match micro evidence on the extensive margin because of heterogeneity in elasticities across subgroups. However, observable heterogeneity in elasticities if anything reinforces the main conclusions drawn above. The heterogeneity in extensive margin responses across groups documented in Table 17 mirrors the heterogeneity observed in extensive margin responses when comparing steady-state behavior across countries with different tax regimes. In particular, individuals near retirement and secondary earners exhibit the largest differences in employment rates across countries with different tax systems (Rogerson and Wallenius 2007, Blundell, Bozio, and Laroque 2013).

In contrast, heterogeneity amplifies the discrepancy between micro and macro estimates of intertemporal substitution elasticities. Employment rates fluctuate substantially over the business cycle even for this subgroup (Clark and Summers 1981, Jaimovich and Siu 2009). This is illustrated by the dashed black series in Figure 25b, which plots detrended employment for males aged 25-55. Fluctuations in employment for prime age males are very similar to those for the population as a whole. However, microeconomic studies clearly show that extensive margin elasticities are near zero for prime-age males. The sharp divergence between micro and macro Frisch elasticities within this subgroup reinforces our conclusion that indivisible labor supply cannot fully account for the fluctuations in aggregate hours over the business cycle.[118]

[118] Fluctuations in wage rates for prime age males are very similar to those for the population as a whole at business cycle frequencies. To illustrate this, we use CPS data on median usual weekly earnings for full time employed wage and salary workers from the Bureau of Labor Statistics (series LEU0252881500) and men aged 25-54 (series LEU0252888100), available from 2000 to 2011. We adjust for inflation using the CPI provided by the BLS aggregated to a quarterly frequency and HP filter the logs of the CPI-adjusted wage series with a smoothing parameter of 1600.

3.7 Conclusion

Indivisible labor is a central feature of many modern macroeconomic models that seek to explain aggregate fluctuations in labor utilization using labor supply. From a qualitative perspective, microeconometric evidence strongly supports the importance of indivisible labor: changes in wage rates clearly induce extensive-margin responses. From a quantitative perspective, observed extensive margin responses are adequate to explain the gap between micro and macro estimates of steady-state elasticities when combined with factors such as frictions. However, extensive margin labor supply responses are not large enough to explain the gap between micro and macro estimates of intertemporal substitution elasticities. Consequently, explanations of the business cycle based on changes in labor supply can only partly explain fluctuations in hours over the business cycle.

One interpretation of our analysis is that it points in favor of recent macro models that feature a cyclical "labor wedge" between the marginal rate of substitution of consumption for leisure and the marginal product of labor. The micro evidence reviewed here is consistent with macro evidence that labor wedges are substantial (Chari et al 2007(Chari, Kehoe, and McGrattan 2007); Shimer, 2009). Our conclusion that labor supply is important but cannot entirely account for fluctuations over the business cycle supports models that combine a labor supply margin with other sources of fluctuations. For instance, Hall (2009) shows that a search-and-matching-generated unemployment margin combined with a labor supply margin can match observed fluctuations in employment rates over the business cycle without requiring large extensive margin labor supply responses.[119] Models that generate unemployment by taking individuals off their labor supply curve in the short run, e.g. due to wage rigidities, are also consistent with our results. While our analysis does not distinguish between alternative explanations of the labor wedge, our estimates could be used to calibrate the labor supply component of models that seek to explain aggregate fluctuations with labor wedges.

Based on our reading of the micro evidence, we recommend calibrating macro models to match Hicksian elasticities of 0.3 on the intensive and 0.25 on the extensive margin and Frisch elasticities

The resulting standard deviation of log real wages around the HP filtered trend is .0122 for the full population and .0123 for prime aged men.

[119] In Hall's model, workers choose both hours and employment based on both standard labor supply factors and the time and effort needed to find a job as in a Diamond-Mortensen-Pissarides model with rigid wages. These forces generate an aggregate hours elasticity of 1.9 even with an intensive Frisch elasticity of 0.7.

of 0.5 on the intensive and 0.25 on the extensive margin.[120,121] These elasticities are consistent with the observed differences in aggregate hours across countries with different tax systems. They also match the relatively small fluctuations in hours conditional on employment over the business cycle. The remaining challenge is to formulate models that fit the large fluctuations in employment rates over the business cycle when calibrated to match an extensive margin labor supply elasticity of 0.25. Even with indivisible labor, models that require a Frisch elasticity of aggregate hours above 1 are inconsistent with micro evidence.

[120] That is, one should choose a reservation wage distribution such that a 10% increase in the net-of-tax wage leads to a 2.5% increase in employment rates. More generally, simulating quasi-experiments such as the tax policy changes analyzed here would be a simple way to evaluate which macro models are consistent with micro data.

[121] We suspect that this estimate is, if anything, biased upward for two reasons: (1) the mean extensive margin elasticity for the population as a whole is less than 0.25 as noted above and (2) publication bias drives micro studies toward reporting higher elasticity estimates (Card and Krueger 1995).

References

Abraham, J. M. and P. H. Hendershott (1996). Bubbles in Metropolitan Housing Markets. *Journal of Housing Research 7*(2), 191–207.

Adelino, M., A. Schoar, and F. Severino (2012). Credit Supply and House Prices: Evidence From Mortgage Market Segementation.

Adjemian, S., H. Bastani, M. Juillard, K. Frederic, F. Mihoubi, G. Perendia, J. Pfeifer, M. Ratto, and S. Villemot (2011). Dynare: Reference Manual, Version 4.

Albrecht, J., A. Anderson, E. Smith, and S. Vroman (2007, May). Opportunistic Matching in the Housing Market. *International Economic Review 48*(2), 641–664.

Albrecht, J., P. A. Gautier, and S. Vroman (2013). Directed Search in the Housing Market.

Alesina, A., E. Glaeser, and B. Sacerdote (2005). Work and Leisure in the U.S. and Europe: Why So DIfferent? In M. Gertler and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2005*, pp. 1–64. Cambridge, MA: MIT Press.

Altonji, J. G. (1986). Intertemporal Substitution in Labor Supply: Evidence From Micro Data. *Journal of Political Economy 94*, S176–S215.

Anenberg, E. (2013). Information Frictions and Housing Market Dynamics.

Anenberg, E. and P. Bayer (2013). Endogenous Sources of Volatility in Housing Markets: The Joint Buyer-Seller Problem.

Anenberg, E. and E. Kung (2014). Estimates of the Size and Source of Price Declines Due to Nearby Foreclosures.

Angeletos, G.-M. and J. La'O (2013). Sentiments. *Econometrica 81*(2), 739–779.

Anglin, P. M., R. Rutherford, and T. Springer (2003). The Trade-off Between the Selling Price of Residential Properties and Time-on-the-Market: The Impact of Price Setting. *Journal of Real Estate Finance and Economics 26*(1), 95–111.

Angrist, J. D. (1991). Grouped-Data Estimation and Testing in Simple Labor-Supply Models. *Journal of Econometrics 47*, 243–266.

Arellano, M. and O. Bover (1995). Another Look at the Instrumental Variable Estimation of Error-Components Models. *Journal of Econometrics 68*, 29–51.

Ashenfelter, O. (1984). Macroeconomic Analyses and Microeconomic Analyses of Labor Supply. *Carnegie-Rochester Conference Series on Public Policy 21*, 117–56.

Auerbach, A. J. (1985). The Theory of Excess Burden and Optimal Taxation. In A. J. Auerbach and M. S. Feldstein (Eds.), *Handbook of Public Economics*, Volume 1, pp. 61–124. Amsterdam: Elsevier.

Bajari, P., S. Chu, and M. Park (2010). An Empirical Model of Subprime Mortgage Default from 2000 to 2007.

Ball, L. and D. Romer (1990). Real Rigidities and the Non-Neutrality of Money. *The Review of Economic Studies 57*(2), 183–203.

Barberis, N., A. Shleifer, and R. Vishny (1998). A Model of Investor Sentiment. *Journal of Financial Economics 49*, 307–343.

Barsky, R., C. L. House, and M. Kimball (2007). Sticky-Price Models and Durable Goods. *American Economic Review 97*(3), 984–998.

Barsky, R., G. Solon, and J. A. Parker (1994). Measuring the Cyclicality of Real Wages: How Important is Composition Bias. *Quarterly Journal of Economics 109*(1), 1–25.

Benabou, R. (1992). Inflation and Efficiency in Search Markets. *The Review of Economic Studies 59*(2), 299–329.

Berkovec, J. A. and J. L. Goodman (1996). Turnover as a Measure of Demand for Existing Homes. *Real Estate Economics 24*(4), 421–440.

Bhutta, N., J. Dokko, and H. Shan (2010). The Depth of Negative Equity and Mortgage Default Decisions During the 2007-2009 Housing Bust.

Bianchi, M., B. R. Gudmundsson, and G. Zoega (2001). Iceland's Natural Experiment in Supply-Side Economics. *American Economic Review 91*(5), 1564–79.

Bils, M., P. J. Klenow, and B. A. Malin (2012). Reset Price Inflation and the Impact of Monetary

Policy Shocks. *American Economic Review 102*(6), 2798–2825.

Blundell, R., A. Bozio, and G. Laroque (2013). Extensive and Intensive Margins of Labour Supply: Working Hours in the U.S., U.K. and France. *Fiscal Studies 34*(1), 1–29.

Blundell, R. and T. E. MaCurdy (1999). Labor Supply: A Review of Alternative Approaches. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, pp. 1559–1695. Amsterdam: Elsevier Science.

Brown, K. M. (2009). The Link Between Pensions and Retirement Timing: Lessons From California Teachers.

Browning, M. (2005). A Working Paper from April 1985: Which Demand Elasticity Do We Know and Which Do We Need to Know for Policy Analysis? *Research in Economics 59*(4), 293–320.

Burnside, C., S. Rebelo, and M. Eichenbaum (2013). Understanding Booms and Busts in Housing Markets.

Calomiris, C. W., S. D. Longhofer, and W. Miles (2008). The Foreclosure-House Price Nexus: Lessons From the 2007-2008 Housing Turmoil.

Calvo, G. (1983). Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics 12*(3), 383–398.

Campbell, J. Y. and J. F. Cocco (2014). A Model of Mortgage Default.

Campbell, J. Y., S. Giglio, and P. Pathak (2011). Forced Sales and House Prices. *American Economic Review 101*(5), 2108–31.

Campbell, J. Y. and N. G. Mankiw (1989). Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence. In O. J. Blanchard and S. Fischer (Eds.), *NBER Macroeconomics Annual 1989*, Volume 4, pp. 185–246. Cambridge, MA: MIT Press.

Caplin, A. and J. Leahy (2011). Trading Frictions and House Price Dynamics. *Journal of Money, Credit and Banking 43*(7), 283–303.

Capozza, D. R., P. H. Hendershott, and C. Mack (2004). An Anatomy of Price Dynamics in

Illiquid Markets: Analysis and Evidence from Local Housing Markets. *Real Estate Economics 32*(1), 1–32.

Card, D. (1990). Labor Supply With a Minimum Hours Threshold. *Carnegie-Rochester Conference Series on Public Policy 33*, 137–168.

Card, D. and D. Hyslop (2005). Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare-Leavers. *Econometrica 73*(6), 1723–70.

Card, D. and A. B. Krueger (1995). Time-Series Minimum-Wage Studies: A Meta-Analysis. *American Economic Review 85*(2), 238–243.

Carrillo, P. E. (2012). An Empirical Stationary Equilibrium Search Model of the Housing Market. *International Economic Review 53*(1), 203–234.

Carrington, W. J. (1996). The Alaskan Labor Market During the Pipeline Era. *Journal of Political Economy 104*(1), 186–218.

Case, K. E. (2008). The Central Role of Home Prices in the Current Financial Crisis: How Will the Market Clear? *Brookings Papers on Economic Activity* (Fall), 161–193.

Case, K. E. and R. J. Shiller (1987). Prices of Single-Family Homes Since 1970: New Indices for Four Cities. *New England Economic Review* (September), 45–56.

Case, K. E. and R. J. Shiller (1989). The Efficiency of the Market for Single-Family Homes. *American Economic Reivew 79*(1), 125–137.

Case, K. E., R. J. Shiller, and A. K. Thompson (2012). What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets. *Brookings Papers on Economic Activity* (Fall), 265–298.

Chang, Y. and S.-B. Kim (2006). From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogenous Agent Macroeconomy. *International Economic Review 47*(1), 1–27.

Chari, V., P. J. Kehoe, and E. R. McGrattan (2000). Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem? *Econometrica 68*(5), 1151–1179.

Chari, V., P. J. Kehoe, and E. R. McGrattan (2007). Business Cycle Accounting. *Econometrica 75*(3), 781–836.

Chatterjee, S. and B. Eyigungor (2011). A Quantitative Analysis of the US Housing and Mortgage Markets and the Foreclosure Crisis.

Chen, Y. and R. W. Rosenthal (1996). Asking Prices as Commitment Devices. *International Economic Review 37*(1), 129–155.

Chetty, R. (2012). Bounds on Elasticities with Optimization Frictions : A Synthesis of Micro and Macro Evidence on Labor Supply. *Econometrica 80*(3), 969–1018.

Chetty, R., J. Friedman, T. Olsen, and L. Pistaferri (2011). The Effect of Adjustment Costs and Institutional Constraints on Labor Supply Elasticities: Evidence from Denmark. *Quarterly Journal of Economics 126*(2), 749–804.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review 101*(3), 471–5.

Chetty, R. and A. Szeidl (2007). Consumption Commitments and Risk Preferences. *The Quarterly Journal of Economics 122*(2), 831–877.

Cho, J.-O. and T. F. Cooley (1994). Employment and Hours Over the Business Cycle. *Journal of Economic Dynamics and Control 18*(2), 411–32.

Cho, J.-O. and R. Rogerson (1988). Family Labor Supply and Aggregate Fluctuations. *Journal of Monetary Economics 21*(2-3), 233–45.

Cho, M. (1996). House Price Dynamics: A Survey of Theoretical and Empirical Issues. *Journal of Housing Research 7*(2), 145–172.

Christiano, L. J. and M. Eichenbaum (1992). Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations. *American Economic Review 82*(3), 430–50.

Clark, K. B. and L. H. Summers (1981). The Demographic Composition of Cyclical Employment Variations. *Journal of Human Resources 16*(1), 61–79.

Clauretie, T. M. and N. Daneshvary (2009). Estimating the House Foreclosure Discount Corrected for Spatial Price Interdependence and Endogeneity of Marketing Time. *Real Estate Economics 37*(1), 43–67.

Coibion, O. and Y. Gorodnichenko (2011). Strategic Interaction Among Heterogeneous Price-Setters in an Estimated DSGE Model. *The Review of Economics and Statistics 93*(3), 920–940.

Cooper, R. and A. John (1988). Coordinating Coordination Failures in Keynesian Models. *Quarterly Journal of Economics 103*(3), 441–463.

Cutler, D. M., J. M. Poterba, and L. H. Summers (1991). Speculative Dynamics. *Review of Economic Studies 58*(3), 529–546.

Davis, S. J. and M. Henrekson (2005). Tax Effects on Work Activity, Industry Mix and Shadow Economy Size: Evidence from Rich-Country Comparisons. In R. G. Salvador, A. Lamo, B. Petrongolo, M. Ward, and E. Wasmer (Eds.), *Labour Supply and Incentives to Work in Europe*, pp. 44–104. Northampton, MA: Edward Elgar Press.

de Wit, E. R. and B. van der Klaauw (2013). Asymmetric Information and List-Price Reductions in the Housing Market. *Regional Science and Urban Economics 43*(3), 507–520.

Deng, Y., J. M. Quigley, and R. Van Order (2000). Mortgage Terminations, Heterogeneity, and the Exercise of Mortgage Options. *Econometrica 68*(2), 275–307.

Devereux, P. (2004). Changes in Relative Wages and Family Labor Supply. *Journal of Human Resources 39*(3), 698–722.

Diaz, A. and B. Jerez (2013). House Prices, Sales, and Time on the Market: A Search-Theoretic Framework. *International Economic Review 54*(3), 837–872.

Domeji, D. and M. Floden (2006). The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased. *Review of Economic Dynamics 9*(2), 242–62.

Duffie, D., N. Garleanu, and L. H. Pedersen (2007). Valuation in Over-the-Counter Markets. *Review of Financial Studies 20*(5), 1865–1900.

Dyrda, S., G. Kaplan, and J.-V. Rios-Rull (2012). Business Cycles and Household Formation: The Micro vs. the Macro Labor Elasticity.

Eckstein, Z. and K. I. Wolpin (1989). Dynamic Labour Force Participation of Married Women and Endogenous Work Experience. *Review of Economic Studies 56*(3), 375–390.

Eissa, N. and H. Hoynes (2004). Taxes and the Labor Market Participation of Married Couples: The Earned Income Tax Credit. *Journal of Public Economics 88*(9-10), 1931–58.

Eissa, N. and J. B. Liebman (1996). Labor Supply Response to the Earned Income Tax Credit. *Quarterly Journal of Economics 111*(2), 605–37.

Elul, R., N. S. Souleles, S. Chomsisengphet, D. Glennon, and R. Hunt (2010). What 'Triggers' Mortgage Default?

Fehr, E. and J.-R. Tyran (2005). Individual Irrationality and Aggregate Outcomes. *Journal of Economic Perspectives 19*(4), 43–66.

Ferreira, F., J. Gyourko, and J. Tracy (2010, July). Housing busts and household mobility. *Journal of Urban Economics 68*(1), 34–45.

Foote, C. L., K. Gerardi, and P. S. Willen (2008). Negative Equity and Foreclosure: Theory and Evidence. *Journal of Urban Economics 64*(2), 234–245.

Frazzini, A. (2006). The Disposition Effect and Underreaction to News. *The Journal of Finance 61*(4), 2017–2046.

Fuhrer, J. C. (2011). Inflation Persistence. In B. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics*, Volume 3, Chapter 9, pp. 423–486. Elsevier.

Fuster, A., D. Laibson, and B. Mendel (2010). Natural Expectations and Macroeconomic Fluctuations. *Journal of Economic Perspectives 24*(4), 67–84.

Gali, J. and M. Gertler (1999). Inflation Dynamics: A Structural Econometric Analysis. *Journal of Monetary Economics 44*, 195–222.

Genesove, D. and L. Han (2012). Search and Matching in the Housing Market. *Journal of Urban Economics 72*(1), 31–45.

Genesove, D. and C. Mayer (1997). Equity and Time to Sale in the Real Estate Market. *American Economic Reivew 87*(3), 255–269.

Genesove, D. and C. Mayer (2001). Loss Aversion and Seller Behavior: Evidence From the Housing Market. *Quarterly Journal of Economics 116*(4), 1233–1260.

Gerardi, K., P. S. Willen, E. Rosenblatt, and V. W. Yao (2013). Foreclosure Externalities: New Evidence.

Glaeser, E. L., J. Gyourko, E. Morales, and C. G. Nathanson (2013). Housing Dynamics.

Golosov, M. and R. E. Lucas (2008). Menu Costs and Phillips Curves. *Journal of Political Economy 115*(2), 171–199.

Gopinath, G. and O. Itskhoki (2010). In Search of Real Rigidities. In D. Acemoglu and M. Woodford (Eds.), *NBER Macroeconomics Annual*, Volume 25, pp. 261–310. Chicago: University of Chicago Press.

Gourio, F. and P.-A. Noual (2009). The Marginal Worker and the Aggregate Elasticity of Labor Supply.

Graversen, E. K. (1998). *Labor Supply and Work Incentives*. Ph. D. thesis, University of Aarhus.

Gruber, J. and D. A. Wise (1999). Introduction and Summary. In J. Gruber and D. A. Wise (Eds.), *Social Security and Retirement Around the World*. Chicago: University of Chicago Press.

Guiso, L., P. Sapienza, and L. Zingales (2013, August). The Determinants of Attitudes Toward Strategic Default on Mortgages. *The Journal of Finance 68*(4), 1473–1515.

Guren, A. and T. McQuade (2013). How Do Foreclosures Exacerbate Housing Downturns?

Hall, R. E. (2009). Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor. *Journal of Political Economy 117*(2), 281–323.

Haltiwanger, J. and M. Waldman (1989). Limited Rationality and Strategic Complements: the Implications for Macroeconomics. *The Quarterly Journal of Economics 104*(3), 463–483.

Han, L. and W. C. Strange (2012). What is the Role of the Asking Price for a House.

Hansen, G. D. (1985). Indivisible Labor and the Business Cycle. *Journal of Monetary Eco-*

nomics 16(3), 309–327.

Haurin, D. R. and H. L. Gill (2002). The Impact of Transaction Costs and the Expected Length of Stay on Homeownership. *Journal of Urban Economics 51*, 563–584.

Haurin, D. R., J. L. Haurin, T. Nadauld, and A. Sanders (2010, December). List Prices, Sale Prices and Marketing Time: An Application to U.S. Housing Markets. *Real Estate Economics 38*(4), 659–685.

Hausman, J. A. (1985). Taxes and Labor Supply. In A. J. Auerbach and M. S. Feldstein (Eds.), *Handbook of Public Economics*, pp. 213–63. Amsterdam.

Hausman, J. A., H. Ichimura, W. K. Newey, and J. L. Powell (1991). Identification and Estimation of Polynomial Errors-In-Variables Models. *Journal of Econometrics 50*(3), 273–295.

Head, A., H. Lloyd-Ellis, and H. Sun (2014). Search, Liquidity, and the Dynamics of House Prices and Construction. *American Economic Review 104*(4), 1172–1210.

Heckman, J. J. (1984). Comments on the Ashenfleter and Kydland Papers. *Carnegie-Rochester Conference Series on Public Policy 21*, 209–24.

Heckman, J. J. (1993). What Has Been Learned About Labor Supply in the Past Twenty Years? *American Economic Review 83*(2), 116–21.

Heckman, J. J. and M. R. Killingsworth (1986). Female Labor Supply: A Survey. In O. Ashenfelter and R. Layard (Eds.), *Handbook of Labor Economics*, pp. 103–204. Amsterdam: Elsevier.

Holtz-Eakin, D., D. Joulfaian, and H. S. Rosen (1993). The Carnegie Conjecture: Some Empirical Evidence. *Quarterly Journal of Economics 108*(2), 413–35.

Hong, H. and J. C. Stein (1999). A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. *The Journal of Finance 54*(6), 2143–2184.

Iacoviello, M. (2005). House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle. *American Economic Reivew 95*(3), 739–764.

Imai, S. and M. Keane (2004). Intertemporal Labor Supply and Human Capital Accumulation. *International Economic Review 45*(2), 601–42.

Imbens, G. W., D. B. Rubin, and B. Sacerdote (2001). Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence From a Survey of Lottery Players. *American Economic Review 91*(4), 778–94.

Jaimovich, N. and H. E. Siu (2009). The Young, the Old, and the Restless: Demographics and Business Cycle Volatility. *American Economic Reivew 99*(3), 804–26.

Juhn, C., K. M. Murphy, and R. H. Topel (1991). Why Has the Natural Rate of Unemployment Increased Over Time? *Brookings Papers on Economic Activity* (2), 75–142.

Kang, H. B. and J. Gardner (1989). Selling Price and Marketing Time in the Residential Real Estate Market. *The Journal of Real Estate Research 4*(1), 21–35.

Keane, M. P. (2010). Labor Supply and Taxes: A Survey.

Keane, M. P. and R. Rogerson (2010). Reconciling Micro and Macro Labor Supply Elasticities.

Kim, J., S. Kim, E. Schaumburg, and C. A. Sims (2008). Calculating and Using Second-Order Accurate Solutions of Discrete Time Dynamic Equilibrium Models. *Journal of Economic Dynamics and Control 32*(11), 3397–3414.

Kimball, M. (1995). The Quantitative Analytics of the Basic Neomonetarist Model. *Journal of Money, Credit and Banking 27*(4), 1241–1277.

Kimmel, J. and T. J. Kniesner (1998). New Evidence on Labor Supply: Employment Versus Hours Elasticities by Sex and Marital Status. *Journal of Monetary Economics 42*(2), 289–301.

King, R. G. and S. T. Rebelo (1999). Resuscitating Real Business Cycles. In J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, pp. 927–1007. Amsterdam: Elsevier.

Kiyotaki, N. and J. Moore (1997). Credit Cycles. *Journal of Political Economy 105*(2), 211–248.

Klenow, P. J. and J. L. Willis (2006). Real Rigidities and Nominal Price Changes.

Knight, J. R. (2002). Listing Price, Time on Market, and Ultimate Selling Price: Causes and Effects of Listing Price Changes. *Real Estate Economics 30*(2), 213–237.

Korteweg, A. and M. Sorensen (2013). Estimating Loan-to-Value Distributions.

Krainer, J. (2001). A Theory of Liquidity in Residential Real Estate Markets. *Journal of Urban Economics 49*(1), 32–53.

Kudoh, N. (2013). Price Competition with Search Frictions.

Lamont, O. and J. C. Stein (1999). Leverage and House-Price Dynamics in U.S. Cities. *RAND Journal of Economics 30*(3), 498–514.

Landvoigt, T., M. Schneider, and M. Piazzesi (2013). The Housing Market(s) of San Diego.

Lazear, E. P. (2010). Why Do Inventories Rise When Demand Falls in Housing and Other Markets?

Leamer, E. E. (2007). Housing is the Business Cycle. In *Housing, Housing Finance and Monetary Policy, A Symposium*, pp. 149–233. Federal Reserve of Kansas City.

Lester, B., L. Visschers, and R. Wolthoff (2013). Competing With Asking Prices.

Levin, A. and T. Yun (2009). Reconsidering the Microeconomic Foundations of Price-Setting Behavior.

Levitt, S. D. and C. Syverson (2008). Market Distortions When Agents Are Better Informed: The Value of Information in Real Estate Transactions. *The Review of Economics and Statistics 90*(4), 599–611.

Liebman, J. B. and E. Saez (2006). Earnings Responses to Increases in Payroll Taxes.

Lin, W., P. K. Robins, D. Card, K. Harknett, and S. Lui-Gurr (1998). *When Financial Incentives Encourage Work: Complete 18-Month Findings From Self-Sufficiency Project*. Ottawa: Social Research and Demonstration Corporation.

Ljungqvist, L. and T. J. Sargent (2006). Do Taxes Explain European Employment? Indivisible Labor, Human Capital, Lotteries and Savings. In D. Acemoglu, K. Rogoff, and M. Woodford (Eds.), *NBER Macroeconomics Annual 2006*, pp. 181–246. Cambridge, MA: MIT Press.

Ljungqvist, L. and T. J. Sargent (2011). A Labor Supply Elasticity Accord? *American Economic*

Review 101(3), 478–91.

Lorenzoni, G. (2009). A Theory of Demand Shocks. *American Economic Review 99*(5), 2050–2084.

Low, H. (2005). Self-Insurance in a Life-Cycle Model of Labor Supply and Savings. *Review of Economic Dynamics 8*(4), 945–75.

Lucas, R. E. (1972). Expectations and the Neutrality of Money. *Journal of Economic Theory 4*, 103–124.

MaCurdy, T. E. (1981). An Empirical Model of Labor Supply in a Life-Cycle Setting. *Journal of Political Economy 89*(6), 1059–1085.

Malpezzi, S. (1999). A Simple Error Correction Model of House Prices. *Journal of Housing Economics 8*(1), 27–62.

Manoli, D. and A. Weber (2011). Nonparametric Evidence on the Effects of Retirement Benefits on Labor Force Participation Decisions.

Meen, G. (2002). The Time-Series Behavior of House Prices: A Transatlantic Divide? *Journal of Housing Economics 11*(1), 1–23.

Meese, R. A. and N. Wallace (1997). The Construction of Residential Housing Price Indices: A Comparison of Repeat-Sales, Hedonic-Regression, and Hybrid Approaches. *Journal of Real Estate Finance and Economics 14*, 51–73.

Meghir, C. and D. Phillips (2010). Labor Supply and Taxes. In T. M. Review (Ed.), *Dimensions of Tax Design*, pp. 202–74. Oxford: Oxford University Press.

Meyer, B. D. (2010). The Effects of the Earned Income Tax Credit and Recent Reforms. In J. M. Poterba (Ed.), *Tax Policy and the Economy*, pp. 153–80. Cambridge, MA: MIT Press.

Meyer, B. D. and D. T. Rosenbaum (2000). Making Single Mothers Work: Recent Tax and Welfare Policy and Its Effects. *National Tax Journal 53*(4), 1027–59.

Meyer, B. D. and D. T. Rosenbaum (2001). Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers. *The Quarterly Journal of Economics 116*(3), 1063–1114.

Mian, A. and A. Sufi (2011). House Prices , Home Equity âĂŞ Based Borrowing , and the U.S. Household Leverage Crisis. *American Economic Review 101*(5), 2132–2156.

Mian, A., A. Sufi, and F. Trebbi (2014). Foreclosures, House Prices, and the Real Economy.

Molloy, R. and H. Shan. The Post-Foreclosure Experience of U.S. Households. (Forthcoming).

Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen (2012, May). Time Series Momentum. *Journal of Financial Economics 104*(2), 228–250.

Mulligan, C. (2001). Aggregate Implications of Indivisible Labor. *Advances in Macroeconomics 1*(1).

Murphy, K. M. and R. H. Topel (1985). Estimation and Inference in Two-Step Econometric Models. *Journal of Business and Economic Statistics 3*(4), 370–379.

Nakamura, E. and D. Zerom (2010). Accounting for Incomplete Pass-Through. *Review of Economic Studies 77*, 1192–1230.

Ngai, L. R. and S. Tenreyro (2013). Hot and Cold Seasons in the Housing Market.

Nickell, S. (2003). Employment and Taxes.

Novy-Marx, R. (2009). Hot and Cold Markets. *Real Estate Economics 37*(1), 1–22.

Ortalo-Magné, F. and S. Rady (2006). Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints. *Review of Economic Studies 73*, 459–485.

Peach, R. (1983). *An Econometric Model of Investment in Single Family Homes*. Ph. D. thesis, University of Maryland.

Pencavel, J. (1986). Labor Supply of Men: A Survey. In O. Ashenfelter and R. Layard (Eds.), *Handbook of Labor Economics*, pp. 3–102. Amsterdam: Elsevier.

Phelps, E. S. (1983). The Trouble With 'Rational Expectations' and the Problem of Inflation Stabilization. In R. Frydman and E. S. Phelps (Eds.), *Individual Forecasting and Aggregate Outcomes*. Cambridge University Press.

Piazzesi, M. and M. Schneider (2009). Momentum Traders in the Housing Market: Survey Evidence and a Search Model. *American Economic Review 99*(2), 406–411.

Pistaferri, L. (2003). Anticipated and Unanticipated Wages Changes, Wage Risk, and Intertemporal Labor Supply. *Journal of Labor Economics 21*(3), 729–54.

Poterba, J. and T. Sinai (2008). Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income. *American Economic Review 98*(2), 84–89.

Prescott, E. C. (2004). Why Do Americans Work So Much More Than Europeans? *Federal Reserve Bank of Minneapolis Quarterly Review 28*(1), 2–13.

Prescott, E. C., R. Rogerson, and J. Wallenius (2009, January). Lifetime Aggregate Labor Supply With Endogenous Workweek Length. *Review of Economic Dynamics 12*(1), 23–36.

Rogerson, R. (1988). Indivisible Labor, Lotteries and Equilibrium. *Journal of Monetary Economics 21*(1), 3–16.

Rogerson, R. and J. Wallenius (2007). Micro and Macro Elasticities in a Life Cycle Model with Taxes.

Rogerson, R. and J. Wallenius (2009). Micro and Macro Elasticities in a Life Cycle Model With Taxes. *Journal of Economic Theory 144*(6), 2277–2292.

Saez, E., J. B. Slemrod, and S. H. Giertz (2012). The Elasticity of Taxable Income With Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature 50*(1), 3–50.

Saiz, A. (2010). The Geographic Determinants of Housing Supply Elasticity. *Quarterly Journal of Economics 125*(3), 1253–1296.

Shimer, R. (2005). The Cyclical Behavior of Equilibrium Unemployment and Vacancies. *American Economic Review 95*(1), 25–49.

Shimer, R. (2009). Convergence in Macroeconomics: The Labor Wedge. *American Economic Journal: Macroeconomics 1*(1), 280–297.

Shimer, R. and I. Werning (2007). Reservation Wages and Unemployment Insurance. *Quarterly*

Journal of Economics 122(3), 1145–1185.

Sinai, T. and N. S. Souleles (2009). Can Owning a Home Hedge the Risk of Moving?

Smets, F. and R. Wouters (2007). Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach. *American Economic Review 97*(3), 586–606.

Springer, T. (1996). Single-Family Housing Transactions: Seller Motivations, Price, and Marketing Time. *Journal of Real Estate Finance and Economics 13*, 237–254.

Stein, J. C. (1995). Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects. *Quarterly Journal of Economics 110*(2), 379–406.

Stiglitz, J. (1979). Equilibrium in Product Markets with Imperfect Information. *American Economic Review 69*(2), 339–345.

Taylor, C. R. (1999). Time-on-the-Market as a Sign of Quality. *Review of Economic Studies 66*(3), 555–578.

Taylor, J. B. (1980). Aggregate Dynamics and Staggered Contracts. *Journal of Political Economy 88*(1), 1–23.

Theologides, S. (2010). Servicing REO Properties : The Servicer's Role and Incentives. In *REO & Vacant Properties: Strategies for Neighborhood Stabilization*, pp. 77–85. Federal Reserve Bank of Boston, Federal Reserve Bank of Cleveland, and Federal Reserve Board.

Trabandt, M. and H. Uhlig (2011). How Far Are We From the Slippery Slope? The Laffer Curve Revisited. *Journal of Monetary Economics 58*(4), 305–327.

Wheaton, W. C. (1990). Vacancy, Search, and Prices in a Housing Market Matching Model. *Journal of Political Economy 98*(6), 1270–1292.

Wheaton, W. C. and N. J. Lee (2009). The Co-Movement of Housing Sales and Housing Prices: Empirics and Theory.

Williams, J. T. (1995). Pricing Real Assets With Costly Search. *Review of Financial Studies 8*(1), 55–90.

Woglom, G. (1982). Underemployment Equilibrium With Rational Expectations. *The Quarterly Journal of Economics 97*(1), 89–107.

Woodford, M. (2003). Imperfect Common Knowledge and the Effects of Monetary Policy. In P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford (Eds.), *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton University Press.

Yavas, A. and S. Yang (1995). The Strategic Role of Listing Price in Marketing Real Estate: Theory and Evidence. *Real Estate Economics 23*(3), 347–368.

Ziliak, J. P. and T. J. Kniesner (1999). Estimating Life Cycle Labor Supply Tax Effects. *Journal of Political Economy 107*(2), 326–59.

Zillow (2010). Price Differences Between Foreclosures and Non-Foreclosures.

A Chapter 1 Appendix

A.1 Data

A.1.1 Time Series Data

National and Regional Data In the main text, three data series are used for price, volume and inventory:

- The CoreLogic national repeat-sales house price index. This is an arithmetic interval-weighted house price index from January 1976 to August 2013. The monthly index is averaged at a quarterly frequency and adjusted for inflation using the Consumer Price Index, BLS series CUUR0000SA0.
- The National Association of Realtors' series of sales of existing single-family homes at a seasonally-adjusted annual rate. The data is monthly for the whole nation from January 1968 to January 2013 and available on request from the NAR. The monthly data is averaged at a quarterly frequency.
- Homes listed for sale comes from vacant homes listed for sale from the Census Housing Vacancy Survey, quarterly from Q1 1968 to Q4 2012. This is divided by the NAR sales volume series to create months of supply.

Other price and inventory measures are used in Appendix A.2. The price measures include:

- A median sales price index for existing single-family homes. The data is monthly for the whole nation from January 1968 to January 2013 and available on request from the National Association of Realtors.
- The quarterly national "expanded purchase-only" HPIs that only includes purchases and supplements the FHFA's database from the GSEs with deeds data from DataQuick from Q1 1991 to Q4 2012. This is an interval-weighted geometric repeat-sales index.
- The monthly Case-Shiller Composite Ten from January 1987 to January 2013. This is an interval-weighted arithmetic repeat-sales index.

- A median sales price index for all sales (existing and new homes) from CoreLogic from January 1976 to August 2013.

The additional inventory measure is the National Association of Realtors' series on inventory and months of supply of existing single-family homes. The data is monthly for the whole nation from June 1982 to February 2013 and is available on request.

For annual AR(1) regressions, I use non-seasonally-adjusted data. Because the volume series comes seasonally adjusted, for any analysis that includes sales volume, I use the data provider's seasonal adjustment if available and otherwise seasonally adjust the data using the Census Bureau's X-12 ARIMA software using a multiplicative seasonal factor.

City-Level Data I create two city-level data sets. The first consists of local repeat-sales price indices for 103 CBSA divisions from CoreLogic. These CBSAs divisions include all CBSAs divisions that are part of the 100 largest CBSAs which have data from at least 1995 onwards. Most of these CBSAs have data starting in 1976. See Table 19 for the full list of CBSAs and years. This data is used for the annual AR(1) regression coefficient histogram in Figure 1 and is adjusted for inflation using the CPI.

The second city-level data set is used for the panel VAR and several cross-city comparisons in Appendix A.2. It combines the same CoreLogic city-level repeat-sales house price indices with transaction volume data for existing home sales from CoreLogic and months of supply at the MSA level provided by the National Association of Realtors. The CoreLogic and NAR data sets are merged using the principal city of the MSA and CBSA division. The volume series sometimes have discontinuities corresponding to the introduction of an additional county to a CBSA, so I examine each time series and select a starting date for the volume series for each CBSA division equal to the month after the last discontinuity. Similarly, the NAR months of supply measure is occasionally not reported for a given MSA. I drop all prior quarters if there are four continuous quarters of missing data. There are, however, a few interspersed quarters with missing data. The similarity between the panel VAR and a VAR on national data shows that the missing quarters are not driving the results. Each MSA's start quarter and end quarter are the first and last quarters, respectively, for which both volume and inventory data are available, with the inventory data typically being the binding constraint. Finally, I limit the sample to 42 MSAs with at least 50 quarters of both

Table 19: CBSAs in CoreLogic City-Level Price Data Set

CBSA Code	Main City Name	Start	End	32820	Memphis, TN	1984	2013
10420	Akron, OH	1978	2013	33124	Miami, FL	1976	2013
10580	Albany, NY	1992	2013	33340	Milwaukee, WI	1976	2013
10740	Albuquerque, NM	1992	2013	33460	Minneapolis, MN	1976	2013
10900	Allentown, PA	1976	2013	34980	Nashville, TN	1976	2013
12060	Atlanta, GA	1976	2013	35004	Nassau, NY	1976	2013
12420	Austin, TX	1976	2013	35084	Newark, NJ-PA	1976	2013
12540	Bakersfield, CA	1976	2013	35300	New Haven, CT	1985	2013
12580	Baltimore, MD	1976	2013	35380	New Orleans, LA	1976	2013
12940	Baton Rouge, LA	1992	2013	35644	New York, NY	1976	2013
13140	Beaumont, TX	1993	2013	36084	Oakland, CA	1976	2013
13644	Bethesda, MD	1976	2013	36420	Oklahoma City, OK	1976	2013
13820	Birmingham, AL	1976	2013	36540	Omaha, NE	1990	2013
14484	Boston, MA	1976	2013	36740	Orlando, FL	1976	2013
14860	Bridgeport, CT	1976	2013	37100	Ventura, CA	1976	2013
15380	Buffalo, NY	1991	2013	37764	Peabody, MA	1976	2013
15764	Cambridge, MA	1976	2013	37964	Philadelphia, PA	1976	2013
15804	Camden, NJ	1976	2013	38060	Phoenix, AZ	1976	2013
16700	Charleston, SC	1976	2013	38300	Pittsburgh, PA	1976	2013
16740	Charlotte, NC	1976	2013	38900	Portland, OR	1976	2013
16974	Chicago, IL	1976	2013	39100	Poughkeepsie, NY	1976	2013
17140	Cincinnati, OH	1976	2013	39300	Providence, RI	1976	2013
17460	Cleveland, OH	1976	2013	39580	Raleigh, NC	1976	2013
17820	Colorado Springs, CO	1976	2013	40060	Richmond, VA	1976	2013
17900	Columbia, SC	1977	2013	40140	Riverside, CA	1976	2013
18140	Columbus, OH	1976	2013	40380	Rochester, NY	1991	2013
19124	Dallas, TX	1977	2013	40484	Rockingham County, NH	1990	2013
19380	Dayton, OH	1976	2013	40900	Sacramento, CA	1976	2013
19740	Denver, CO	1976	2013	41180	St. Louis, MO	1978	2013
19804	Detroit, MI	1989	2013	41620	Salt Lake City, UT	1992	2013
20764	Edison, NJ	1976	2013	41700	San Antonio, TX	1991	2013
21340	El Paso, TX	1977	2013	41740	San Diego, CA	1976	2013
22744	Fort Lauderdale, FL	1976	2013	41884	San Francisco, CA	1976	2013
23104	Fort Worth, TX	1984	2013	41940	San Jose, CA	1976	2013
23420	Fresno, CA	1976	2013	42044	Santa Ana, CA	1976	2013
23844	Gary, IN	1992	2013	42644	Seattle, WA	1976	2013
24340	Grand Rapids, MI	1992	2013	44140	Springfield, MA	1976	2013
24660	Greensboro, NC	1976	2013	44700	Stockton, CA	1976	2013
24860	Greenville, SC	1976	2013	45060	Syracuse, NY	1992	2013
25540	Hartford, CT	1976	2013	45104	Tacoma, WA	1977	2013
26180	Honolulu, HI	1976	2013	45300	Tampa, FL	1976	2013
26420	Houston, TX	1982	2013	45780	Toledo, OH	1976	2013
26900	Indianapolis, IN	1991	2013	45820	Topeka, KS	1985	2013
27260	Jacksonville, FL	1976	2013	46060	Tucson, AZ	1976	2013
28140	Kansas City, MO	1985	2013	46140	Tulsa, OK	1981	2013
28940	Knoxville, TN	1977	2013	47260	Virginia Beach, VA	1976	2013
29404	Lake County, IL	1982	2013	47644	Warren, MI	1976	2013
29820	Las Vegas, NV	1983	2013	47894	Washington, DC	1976	2013
30780	Little Rock, AR	1985	2013	48424	West Palm Beach, FL	1976	2013
31084	Los Angeles, CA	1976	2013	48620	Wichita, KS	1986	2013
31140	Louisville, KY	1987	2013	48864	Wilmington, DE	1976	2013
32580	McAllen, TX	1992	2013	49340	Worcester, MA	1976	2013

Table 20: MSAs in Merged Price-Inventory-Volume Panel

Principal City	First Date	Last Date	Obs	Principal City	First Date	Last Date	Obs
Akron, OH	1990, Q1	2013, Q1	87	Miami, FL	1993, Q1	2013, Q1	69
Allentown, PA	1999, Q1	2013, Q1	50	Milwaukee, WI	1998, Q1	2013, Q1	61
Atlanta, GA	1999, Q1	2012, Q3	55	Nashville, TN	1999, Q3	2013, Q1	55
Austin, TX	1994, Q1	2013, Q1	66	New Brunswick, NJ	1999, Q1	2013, Q1	57
Baltimore, MD	1999, Q1	2013, Q1	57	New York, NY	1999, Q1	2013, Q1	57
Charleston, SC	1997, Q1	2013, Q1	65	Newark, NJ	1998, Q1	2013, Q1	61
Chicago, IL	1999, Q1	2013, Q1	57	Oklahoma City, OK	1990, Q1	2013, Q1	79
Cincinnati, OH	1990, Q1	2006, Q4	51	Omaha, NE-IA	2000, Q3	2013, Q1	51
Columbia, SC	1995, Q1	2013, Q1	61	Phoenix, AZ	1993, Q2	2012, Q2	74
Columbus, OH	1995, Q2	2013, Q1	69	Portland, OR	1994, Q1	2011, Q3	66
Dallas, TX	2000, Q1	2013, Q1	53	Providence, RI	1994, Q1	2013, Q1	71
Denver, CO	1999, Q1	2013, Q1	57	Raleigh, NC	1991, Q2	2008, Q2	55
Greenville, SC	1994, Q1	2013, Q1	60	Richmond, VA	1990, Q1	2009, Q2	57
Honolulu, HI	1999, Q1	2013, Q1	57	San Antonio, TX	1998, Q2	2013, Q1	60
Houston, TX	1999, Q3	2013, Q1	55	San Diego CA	1997, Q1	2013, Q1	65
Kansas City, MO	1998, Q3	2012, Q3	57	San Francisco, CA	1992, Q2	2009, Q4	58
Knoxville, TN	1998, Q1	2013, Q1	57	Santa Ana, CA	1999, Q1	2013, Q1	57
Las Vegas, NV	1992, Q2	2013, Q1	63	St. Louis, MO	1996, Q2	2013, Q1	64
Little Rock, AR	1998, Q3	2013, Q1	59	Tampa, FL	1999, Q1	2013, Q1	57
Los Angeles, CA	1993, Q3	2013, Q1	79	Tulsa, OK	1999, Q1	2013, Q1	57
Memphis, TN	1994, Q2	2013, Q1	76	Washington, DC	1993, Q2	2013, Q1	72

inventory and volume data. Table 20 summarizes the full list of MSAs and years in the data set.

A.1.2 Micro Data

The matched listings-transactions micro data covers the San Francisco Bay, San Diego, and Los Angeles metropolitan areas. The San Francisco Bay sample includes Alameda, Contra Costa, Marin, San Benito, San Francisco, San Mateo, and Santa Clara counties. The Los Angeles sample includes Los Angeles and Orange counties. The San Diego sample only includes San Diego County. The data from DataQuick run from January 1988 to August 2013. The Altos data run from October 2007 to May 2013. I limit my analysis to April 2008 to February 2013, as described in footnote 14.

DataQuick Characteristic and History Data Construction The DataQuick data is provided in separate assessor and history files. The assessor file contains house characteristics from the property assessment and a unique property ID for every parcel in a county. The history file contains records of all deed transfers, with each transfer matched to a property ID. Several steps are used to clean the data.

First, both data files are formatted and sorted into county level data files. For a very small number of properties, data with a typo is replaced as missing.

Second, some transactions appear to be duplicates. Duplicate values are categorized and combined into one observation if possible. I drop cases where there are more than ten duplicates, as this is usually a developer selling off many lots individually after splitting them. Otherwise, I pick the sale with the highest price, or, if as a tiebreaker, the highest loan value at origination. In practice, this affects very few observations.

Third, problematic observations are identified. In particular, transfers between family members are identified and dropped based on a DataQuick transfer flag and a comparison buyer and seller names. Sales with prices that are less than or equal to one dollar are also counted as transfers. Partial consideration sales, partial sales, group sales, and splits are also dropped, as are deed transfers that are part of the foreclosure process but not actually transactions. Transactions that appear to be corrections or with implausible origination loan to value ratios are also flagged and dropped. Properties with implausible characteristics (<10 square feet, < 1 bedroom, $< 1/2$ bathroom, implausible year built) have the implausible characteristic replaced as a missing value.

From the final data set, I only use resale transactions (as opposed to new construction or subdivisions) of single-family homes, both of which are categorized by DataQuick.

Altos Research Listings Data Construction and Match to DataQuick The Altos research data contains address, MLS identifier, house characteristics, list price, and date for every week-listing. Altos generously provided me access to an address hash that was used to parse the address fields in the DataQuick assessor data and Altos data and to create a matching hash for each. Hashes were only used that appeared in both data files, and hashes that matched to multiple DataQuick properties were dropped.

After formatting the Altos data, I match the Altos data to the DataQuick property IDs. I first use the address hash, applying the matched property ID to every listing with the same MLS identifier (all listings with the same MLS ID are the same property, and if they do not all match it is because some weeks the property has the address listed differently, for instance "street" is included in some weeks but not others). Second, I match listings not matched by the address hash by repeatedly matching on various combinations of address fields and discarding possible matches

when there is not a unique property in the DataQuick data for a particular combination of fields, which prevents cases where there are two properties that would match from being counted as a match. I determined the combinations of address fields on which to match based on an inspection of the unmatched observations, most of which occur when the listing in the MLS data does not include the exact wording of the DataQuick record (*e.g.*, missing "street"). The fields typically include ZIP, street name, and street number and different combinations of unit number, street direction, and street suffix. In some cases I match to the first few digits of street number or the first word of a street name. I finally assign any unmatched observations with the same MLS ID as a matched observation or the same address hash, square feet, year built, ZIP code, and city as a matched observation the property ID of the matched observation. I subsequently work only with matched properties so that I do not inadvertently count a bad match as a withdrawal.

The observations that are not matched to a DataQuick property ID are usually multi-family homes (which I subsequently drop), townhouses with multiple single-family homes at the same address, or listings with typos in the address field.

I use the subset of listings matched to a property ID and combine cases where the same property has multiple MLS identifiers into a contiguous listing to account for de-listings and re-listings of properties, which is a common tactic among real estate agents. In particular, I count a listing as contiguous if the property is re-listed within 13 weeks and there is not a foreclosure between the de-listing and re-listing. I assign each contiguous listing a single identifier, which I use to match to transactions.

In a few cases, a listing matches to several property IDs. I choose the property ID that matches to a transaction or that corresponds to the longest listing period. All results are robust to dropping the small number of properties that match to multiple property IDs.

I finally match all consolidated listings to a transaction. I drop transactions and corresponding listings where there was a previous transaction in the last 90 days, as these tend to be a true transaction followed by several subsequent transfers for legal reasons (*e.g.*, one spouse buys the house and then sells half of it to the other). I first match to a transaction where the date of last listing is in the month of the deed transfer request or in the prior three months. I then match unmatched listings to a transaction where the date of last listing is in the three months after the deed transfer request (if the property was left on the MLS after the request, presumably by

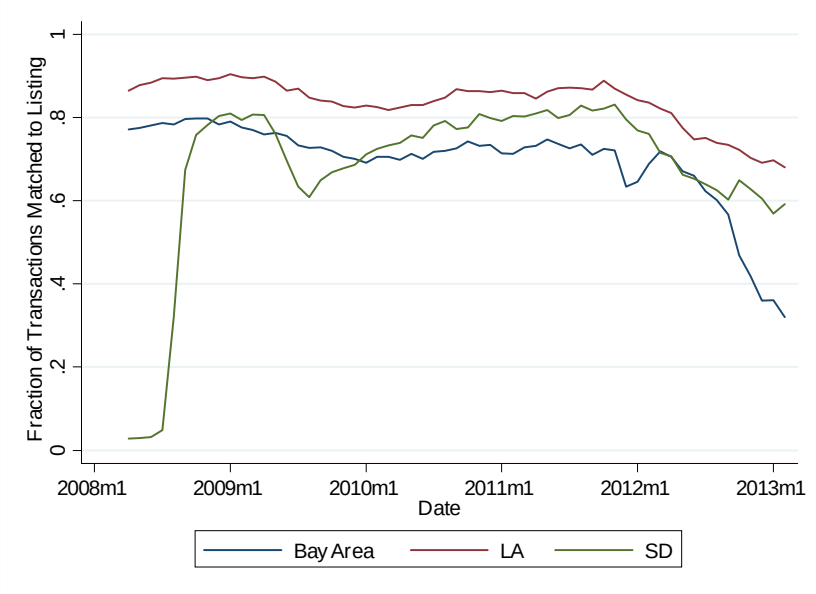

Figure 26: Match Rates by Month of Transaction

accident). I then repeat the process for unmatched listings for four to 12 months prior and four to 12 months subsequent. Most matches have listings within three months of the last listing.

For matched transactions, I generate two measures of whether a house sold within a given time frame. The first, used in the main text, is the time between the date of first listing and the date of filing of the deed transfer request. The second, used in robustness checks in Appendix A.3, is the time between date of first listing and the first of the last listing date or the transfer request.

Figure 26 shows the fraction of all single-family transactions of existing homes for which my data accounts in each of the three metropolitan areas over time. Because the match rates start low in October 2007, I do not start my analysis until April 2008, except in San Diego where almost all listings have no listed address until August 2008. Besides that, the match rates are fairly stable, except for a small dip in San Diego in mid-2009 and early 2012 and a large fall off in the San Francisco Bay area after June 2012. I consequently end the analysis for the San Francisco Bay area at June 2012. Figures 27, 28, and 29 show match rates by ZIP code. One can see that the match rate is consistently high in the core of each metropolitan area and falls off in the outlying areas, such as western San Diego county and Escondido in San Diego, Santa Clarita in Los Angeles, and Brentwood and Pleasanton in the San Francisco Bay area.

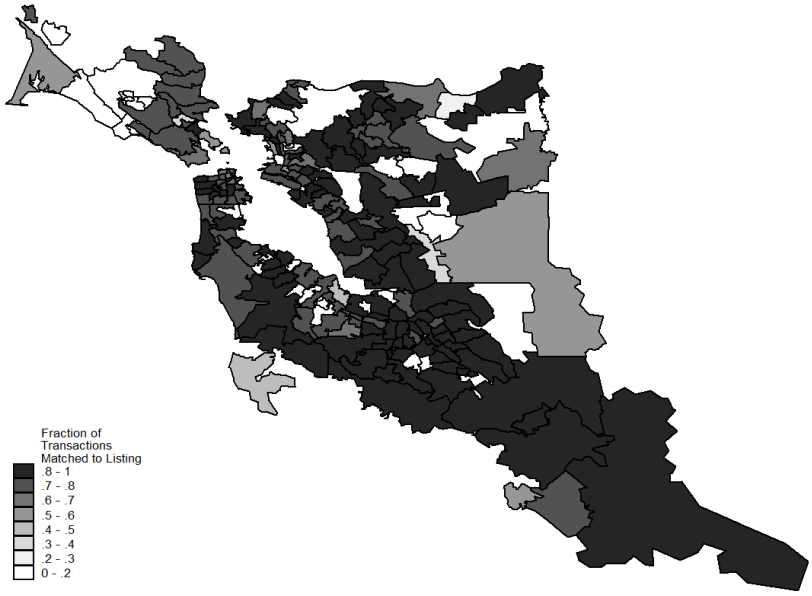

Figure 27: Match Rates by ZIP Code: Bay Area

Figure 28: Match Rates by ZIP Code: Los Angeles

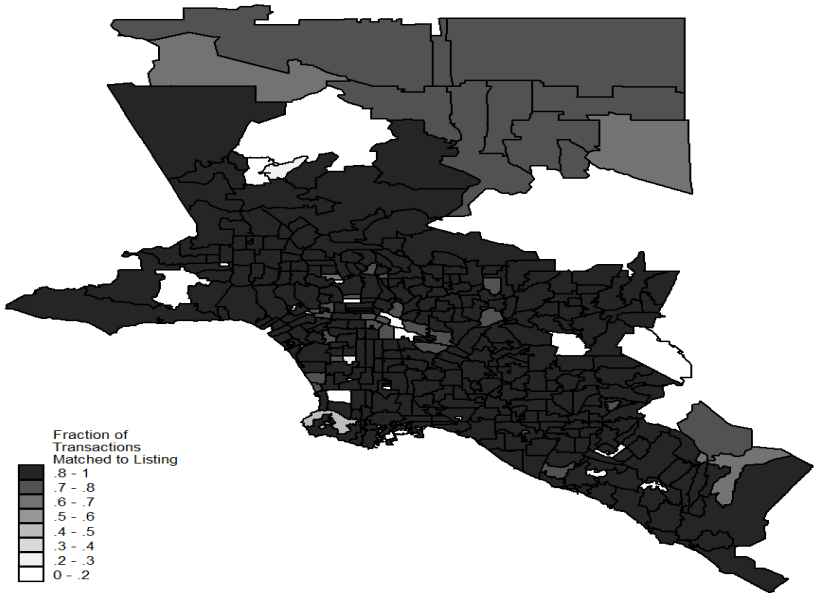

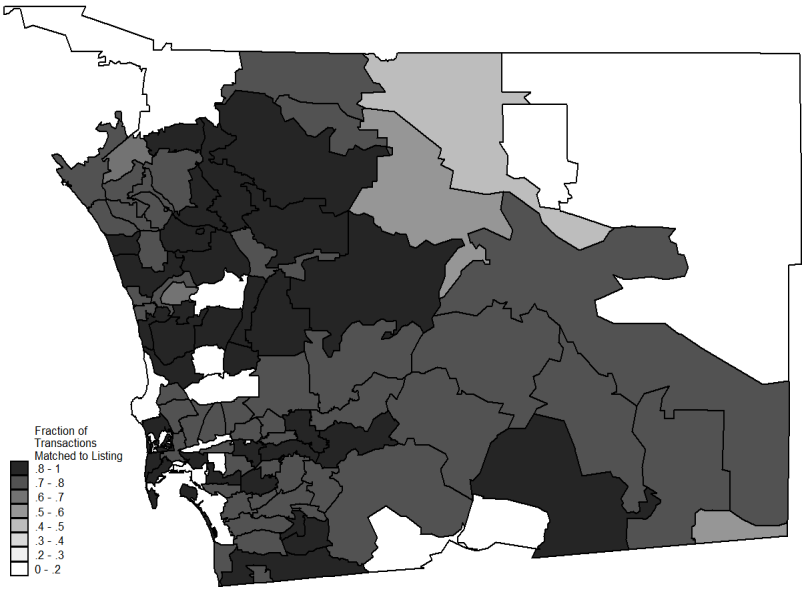

Figure 29: Match Rates by ZIP Code: San Diego

Construction of House Price Indices I construct house price indices largely following Case and Shiller (1989) and follow sample restrictions imposed in the construction of the Case-Shiller and Federal Housing Finance Administration (FHFA) house price indices.

For the repeat sale indices, I drop all non-repeat sales, all sales pairs with less than six months between sales, and all sales pairs where a first stage regression on year dummies shows a property has appreciated by 100 percent more or 100 percent less than the average house in the MSA. I estimate an interval-corrected geometric repeat-sales index at the ZIP code level. This involves estimating a first stage regression:

$$p_{h\ell t} = \xi_{h\ell} + \phi_t + \varepsilon_{h\ell t}, \tag{40}$$

where p is the log price of a house h in location ℓ at time t, $\xi_{h\ell}$ is a sales pair fixed effect, ϕ_t is a time fixed effect, and $\varepsilon_{h\ell t}$ is an error term.

I follow Case and Shiller (1989) by using a GLS interval-weighted estimator to account for the fact that longer time intervals tend to have a larger variance in the error of (40). This is typically implemented by regressing the square of the error term $\varepsilon^2_{h\ell t}$ on a linear (Case-Shiller index) or quadratic (FHFA) function of the time interval between the two sales. The regression coefficients

are then used to construct weights corresponding to $\frac{1}{\sqrt{\hat{\varepsilon}^2_{h\ell t}}}$ where $\hat{\varepsilon}^2_{h\ell t}$ is a predicted value from the interval regression. I find that the variance of the error of (40) is non-monotonic: it is very high for sales that occur quickly, falls to its lowest level for sales that occur approximately three years after the first sale, and then rises slowly over time. This is likely due to flippers who upgrade a house and sell it without the upgrade showing up in the data. Consequently, I follow a non-parametric approach by binning the data into deciles of the time interval between the two sales, calculate the average $\varepsilon^2_{h\ell t}$ for the decile $\bar{\varepsilon}^2_{h\ell t}$, and weight by $\frac{1}{\sqrt{\bar{\varepsilon}^2_{h\ell t}}}$. The results are nearly identical using a linear interval weighting.

$\exp(\phi_t)$ is then a geometric house price index. The resulting indices can be quite noisy. Consequently, I smooth the index using a 3-month moving average, which produced the lowest prediction error of several different window widths. The resulting indices at the MSA level are very comparable to published indices by Case-Shiller, the FHFA, and CoreLogic.

The log predicted value of a house at time t, $\hat{p}_t$, that sold originally at time τ for P_τ is:

$$\hat{p}_t = \log\left(\frac{\exp\left(\hat{\phi}_t\right)}{\exp\left(\hat{\phi}_\tau\right)}P_\tau\right).$$

For the hedonic house price indices, I use all sales and estimate:

$$p_{i\ell t} = \phi_t + \beta X_i + \varepsilon_{i\ell t}, \tag{41}$$

where X_i is a vector of third-order polynomials in four housing characteristics: age, bathrooms, bedrooms, and $\log(\text{square feet})$, all of which are winsorized at the one percent level by county for all properties in a county, not just those that trade. Recall that these characteristics are all recorded as a single snapshot in 2013, so X_i is not time dependent. I do not include a characteristic if over 25 percent of the houses in a given geography are missing data for a particular characteristic. Again $\exp(\phi_t)$ is a house price index, which I smooth using a 3-month moving average. The log predicted price of a house is

$$\hat{p}_{it} = \hat{\beta}X_i + \hat{\phi}_i.$$

For homes that are missing characteristics included in an area's house price index calculation, I

replace the characteristic with its average value in a given ZIP code.

For robustness I calculate both indices for the full sample and a non-distressed sample, where a repeat-sales pair counts as distressed if either sale is an REO sale, a foreclosure auction sale, or a short sale. For my analysis, I use a ZIP code level index, but all results are robust to alternatively using a house price index for all homes within one mile of the centroid of a home's seven-digit ZIP code (roughly a few square blocks). I do not calculate a house price index if the area has fewer than 500 sales since 1988. This rules out about 5% of transactions, typically in low-density areas far from the core of the MSA. For each ZIP code, I calculate the standard deviation of the prediction error of the house price index from 1988 to 2013 and weight most specifications by the reciprocal of the standard deviation.

Construction of the Final Analysis Samples I drop listings that satisfy one of several criteria:

1. If the list price is less than $10,000;
2. If the assessed structure value is less than five percent of the assessed overall value;
3. If the data shows the property was built after the sale date or there has been "significant improvement" since the sale date;
4. If there is a previous sale within 90 days.

Each observation is a listing, regardless of whether it is withdrawn or ends in a transaction. The outcome variable is sold within 13 weeks, where withdrawn listings are counted as not transacting. The price variable is the initial list price. The predicted prices are calculated for the week of first listing by interpolation. The sample is summarized in Table 2 in the main text, and the fraction of the sample accounted for by each MSA and year are summarized in Table 21.

A.1.3 List Prices Relative to Transaction Prices

As mentioned in the main text, the modal house sells at its list price at the time of sale and the average and median house sell within 0.01 log points of their list price. To illustrate this, Figure 30 shows a histogram of the difference between the log list price at sale and the log transaction price

Table 21: Share of Sample Accounted For By Each MSA and Year

Sample	All All	Prior Trans All	All Transactions	Prior Trans Transactions
SF Bay	26.98%	26.68%	28.03%	27.44%
Los Angeles	58.81%	59.47%	57.37%	58.22%
San Diego	14.22%	13.85%	14.59%	14.33%
2008	18.18%	19.87%	16.42%	17.91%
2009	20.70%	21.26%	21.19%	21.81%
2010	23.88%	23.59%	23.48%	23.15%
2011	21.07%	20.36%	21.64%	20.97%
2012	14.86%	13.75%	15.93%	14.93%
2013	1.30%	1.17%	1.34%	1.24%

Notes: Each cell indicates the percentage of each sample accounted for by each MSA (above the line) or by each year of first listing (below the line).

Figure 30: Histogram of the Difference Between Log Transaction Price and Log List Price

Notes: The figure shows a histogram of the difference between log transaction price at the time of sale and log list price for all homes in the San Francisco Bay, Los Angeles, and San Diego areas that were listed between April 2008 and February 2013 that are matched to a transaction and have a previous observed listing. The 1st and 99th percentiles are dropped from the histogram. N = 303,731.

in the Altos-DataQuick merged data. One can see that nearly 18 percent of transactions sell at list price, and the mean of the list price distribution is 0.01 log points below the transaction price.

Table 22 reinforces these findings by showing mean log difference for each of the three MSAs in each year. The mean does not fluctuate by more than 0.03 log points across years and MSAs.

Table 22: Mean Difference Between Log Transaction Price and Log List Price

	SF Bay	Los Angeles	San Diego
2008	-0.007	-0.014	0.003
2009	0.000	-0.008	-0.006
2010	0.001	-0.011	-0.015
2011	-0.014	-0.022	-0.028
2012-3	0.001	-0.015	-0.019

Notes: Each cell shows the mean difference between the log transaction price and log list price in the indicated MSA-year cell. N = 303,731

A.2 Housing Market Facts

A.2.1 Momentum

To assess the robustness of the facts about house price momentum presented in Section 1.2, Table 23 shows several measures of momentum for five different national price indices. The indices are the CoreLogic National repeat-sales house price index discussed in the main text, the Case-Shiller Composite Ten, the FHFA expanded repeat-sales house price index, the National Association of Realtors' national median price for single-family homes, and CoreLogic's national median price for all transactions. The first column shows the coefficient on an AR(1) in log annual price change run at quarterly -frequency as in equation (1).[122] The next two columns show the one- and two-year lagged autocorrelations of the quarterly change in log price. The fourth column shows the quarterly lag in which the autocorrelation of the quarterly change in log price is first negative. The fifth column shows the quarter subsequent to a shock in which the impulse response from an estimated AR(5) estimated in log levels, as in Section 1.2, reaches its peak value. Finally, the sixth column shows the quarterly lag in which the Lo-MacKinlay variance ratio statistic reaches its peak value. This statistic is equal to,

$$V(k) = \frac{var\left(\sum_{t=1}^{t-k+1} r_{t-k+1}\right)/k}{var(r_t)} = \frac{var\left(\log(p_t) - \log(p_{t-k})\right)/k}{var\left(\log(p_t) - \log(p_{t-1})\right)}, \tag{42}$$

[122] Case and Shiller (1989) worry that the same house selling twice may induce correlated errors that generate artificial momentum in regression (1) and use $\Delta p_{t,t-4}$ from one half of their sample and $\Delta p_{t-4,t-8}$ from the other. I have found that this concern is minor with 25 years of administrative data by replicating their split sample approach with my own house price indices estimated from the micro data.

Table 23: The Robustness of Momentum Across Price Measures and Metrics

Price Measure	Annual AR(1) Coefficient	1 Year Lagged Autocorr of Quarterly Δp	2 Year Lagged Autocorr of Quarterly Δp	Lag in Which Quarterly Δp Autocorr is First < 0	Quarter of Peak of AR(5) IRF	Quarter of Peak Value of Lo-MacKinlay Variance Ratio
CoreLogic Repeat Sales HPI, 1976-2013	0.665 (0.081)	0.516	0.199	12	12	19
Case-Shiller Comp 10, 1987-2013	0.67 (0.088)	0.578	0.251	14	11	20
FHFA Expanded HPI, 1991-2013	0.699 (0.089)	0.585	0.344	14	11	18
NAR Median Price, 1968-2013	0.458 (0.103)	0.147	0.062	12	6	16
CoreLogic Median Price, 1976-2013	0.473 (0.082)	0.215	0.046	11	7	16

Notes: Each row shows six measures of momentum for each of the five house price indices, which are detailed in Appendix A.1. The first row shows the AR(1) coefficient for a regression of the annual change in log price on a on-year lag of itself estimated on quarterly data, as in equation (1), with robust standard errors in parenthesis. The second and third columns show the one and two year lagged autocorrelations of the quarterly change in log price. The fourth column shows the quarterly lag in which the autocorrelation of the quarterly change in log price is first negative. The fifth column indicates the quarter in which the impulse response function estimated from an AR(5), as in Section 1.2, reaches its peak. Finally, the last column shows the quarterly lag for which the Lo-MacKinlay variance ratio computed as in equation (42) reaches its peak.

Table 24: Testing For Asmmetry in Momentum

Dependent Variable: Annual Change in Log Price Index at CBSA Level		
Specification	With Interaction	Without Interaction
Coefficient on Year-Lagged	0.614***	0.591***
Annual Change in Log Price	(0.011)	(0.020)
Coefficient on Interaction With		0.045
Positive Lagged Change		(0.031)
CBSA Fixed Effects	Yes	Yes
CBSAs	103	103
N	13,188	13,188

Notes: *** p<0.001. Each column shows a regression of the annual change in log price on a one-year lag of itself and CBSA fixed effects. In column two, the interaction between the lag of annual change in log price with an indicator for whether the lag of the annual change in log price is also included as in equation (43). The regressions are estimated on the panel of 103 CBSAs repeat-sales price indices described in Appendix A.1. Robust standard errors are in parentheses.

where $r_t = \log(p_t) - \log(p_{t-1})$ is the one-period return. If this statistic is equal to one, then there is no momentum, and several papers have used the maximized period of the statistic as a measure of the duration of momentum.

Table 23 shows evidence of significant momentum for all price measures and all measures of momentum. The two median price series exhibit less momentum as the IRFs peak at just under two years and the two-year-lagged autocorrelation is much closer to zero.

Table 24 tests for asymmetry in momentum. Many papers describe prices as being primarily sticky on the downside (*e.g.*, Leamer, 2007; Case, 2008). To assess whether this is the case, I turn to the panel of 103 CBSA repeat-sales price indices described in Appendix A.1, which allows for a more powerful test of asymmetry than using a single national data series. I estimate a quarterly AR(1) regression of the form:

$$\Delta_{t,t-4} \ln p_c = \beta_0 + \beta_1 \Delta_{t-4,t-8} \ln p_c + \beta_2 \Delta_{t-4,t-8} \ln p_c \times 1 \left[\Delta_{t-4,t-8} \ln p_c > 0\right] + \phi_c + \varepsilon, \qquad (43)$$

where c is a city. If momentum is stronger on the downside, the interaction coefficient β_2 should be negative. However, Table 24 shows that the coefficient is insignificant and positive. Thus momentum appears equally strong on the upside and downside when measured using a repeat-sales index.

Across Countries Table 25 shows annual AR(1) regressions as in equation (1) run on quarterly non-inflation-adjusted data for ten countries. The data come from the Bank for International Settlements, which compiles house price indices from central banks and national statistical agencies. The data and details can be found online at http://www.bis.org/statistics/pp.htm. I select ten countries from the BIS database that include at least 15 years of data and have a series for single-family detached homes or all homes. Countries with per-square-foot indices are excluded. With the exception of Norway, which shows no momentum, and the Netherlands, which shows anomalously high momentum, all of the AR(1) coefficients are significant and between 0.2 and 0.6. Price momentum thus appears to show up across countries as well as within the United States and across U.S. metropolitan areas.

Table 25: Momentum Across Countries

Country	AR(1) Coefficient	N	Country	AR(1) Coefficient	N
Australia, 1986-2013	0.217*	100	Netherlands, 1995-2013	0.951***	67
	(0.108)			(0.079)	
Belgium, 1973-2013	0.231**	154	Norway, 1992-2013	-0.042	79
	(0.074)			(0.091)	
Denmark, 1992-2013	0.412***	78	New Zealand, 1979-2013	0.507***	127
	(0.110)			(0.075)	
France, 1996-2013	0.597***	62	Sweden, 1986-2013	0.520***	103
	(0.121)			(0.100)	
Great Britain, 1968-2013	0.467***	173	Switzerland, 1970-2013	0.619***	167
	(0.079)			(0.082)	

Notes: * p < 0.05, ** p<0.01, *** p<0.001. Each row shows the AR(1) coefficient for a regression of the annual change in log price on an annual lag of itself, as in equation (1), estimated on quarterly, non-inflation-adjusted data from the indicated country for the indicated time period. Robust standard errors are in parentheses, and N indicates the number of quarters in the sample. The BIS identifiers and series descriptions are listed for each country. Australia: Q:AU:4:3:0:1:0:0, residential property for all detached houses, eight cities. Belgium Q:BE:0:3:0:0:0:0, residential property all detached houses. Denmark: Q:DK:0:2:0:1:0:0, residential all single-family houses. France: Q:FR:0:1:1:6:0, residential property prices of existing dwellings. Great Britain: Q:GB:0:1:0:1:0:0, residential property prices all dwellings from the Office of National Statistics. Netherlands: Q:NL:0:2:1:1:6:0, residential existing houses. Norway: Q:NO:0:3:0:1:0:0, Residential detached houses. New Zealand: Q:NZ:0:1:0:3:0:0, residential all dwellings. Sweden: Q:SE:0:2:0:1:0:0, owner-occupied detached houses. Switzerland: Q:CH:0:2:0:2:0:0, owner-occupied single-family houses.

A.2.2 Housing Cycle Facts

Relative Volatilities To assess the robustness the relative volatilities of price, volume, and inventory summarized by Fact 2, Table 26 shows the standard deviation of annual log changes for four additional measures of price and two additional measures of inventory discussed in Appendix A.1. The series all have different time coverages, and the standard deviation is calculated for the period over which data is available. The first three rows show various house price indices. With the exception of the Case-Shiller Composite Ten, which is known to be volatile given that it follows ten cities with relatively inelastic housing supplies, the standard deviation of annual log changes are close to the 0.065 figure for the national CoreLogic house price index presented in the main text, although the two median price series are slightly less volatile. The last two rows show two measures of inventory. The first is houses listed for sale rather than months of supply. This is about half as volatile as months of supply because it is not divided by volume. The second is a separate months of supply measure from the NAR, which is only slightly less volatile than the measure in the main

Table 26: Robustness of Relative Volatilities

Measure	$\sigma_{\log x_q - \log x_{q-4}}$
Case-Shiller Composite 10, 1986-2013	0.087
FHFA Expanded, 1991-2013	0.052
NAR Median Price, 1968-2013	0.046
CoreLogic Median Price, 1976-2013	0.061
Census For Sale Inventory, 1968-2013	0.106
NAR Months of Supply, 1982-2013	0.170

Notes: All series are 1976-2013 at a quarterly frequency. The first column shows the standard deviation of annual changes. Data is described in Appendix A.1.

text that combines Census and NAR data. I was not able to find another national volume series with long enough coverage to reliably calculate volatility. Overall, Table 26 supports the conclusion that price is less volatile than inventory and volume, regardless of how each is measured.

Housing Phillips Curve To assess the robustness of the "housing Phillips curve" relationship between price changes and inventory levels (Fact 3), Table 27 shows regression coefficients and R-squareds for regressions of the annual change in log price on log inventory levels, as in equation (26), for five different house price indices and three different measures of inventory. The strong negative relationship is present across all 15 combinations of price and inventory measures.

Two things in particular are of note. First, the relationship is stronger for repeat-sales house price indices, which display more momentum, than it is for median price indices. Second, the middle row shows that the result is robust to measuring inventory as homes listed for sale (adjusted for a linear time trend) instead of as months of supply. The importance of homes listed for sale, the numerator of months of supply, suggests that the price-volume relationship, which affects the denominator of months of supply, is not driving the negative relationship between price changes and inventory levels in the data.

Price, Volume, and Inventory VAR and VEC In the main text, I estimate a panel vector autoregression model on log price, log volume, and log inventory on a panel of 42 cities. The model for the panel VAR is:

$$x_{ct} = \Gamma_0 + \Gamma_1 x_{c,t-1} + \Gamma_2 x_{c,t-2} + \varphi_c + \varepsilon_t$$

Table 27: Robustness of Housing Phillips Curve Relationship

Inventory Measure		CoreLogic National HPI 1976-2013	Case-Shiller Composite Ten 1987-2013	FHFA Expanded Natl HPI 1991-2013	NAR Median Price HPI 1968-2013	CoreLogic Median Price HPI 1976-2013
Months of Supply,	β	-0.140***	-0.232***	-0.147***	-0.089***	-0.095***
1968-2013		(0.015)	(0.022)	(0.010)	(0.011)	(0.014)
From Census & NAR	R^2	0.530	0.630	0.804	0.406	0.283
Homes For Sale	β	-0.283***	-0.385***	-0.273***	-0.148***	-0.235***
1968-2013		(0.035)	(0.054)	(0.029)	(0.021)	(0.030)
From Census	R^2	0.404	0.407	0.619	0.310	0.319
Months of Supply	β	-0.130***	-0.237***	-0.170***	-0.088***	-0.083***
1982-2013		(0.016)	(0.023)	(0.013)	(0.013)	(0.017)
From NAR	R^2	0.367	0.597	0.792	0.283	0.191

Notes: * p < 0.05, ** p<0.01, *** p<0.001. Each cell shows a regression of the annual change in log price on log inventory levels measured at the midpoint of the year over which changes are calculated, as in equation (26). Each column uses a different measure of price, and each row uses a different measure of inventory. For prices, the measures used are the CoreLogic national repeat-sales HPI from 1976-2013 as in the main text, the Case-Shiller Composite 10 repeat-sales HPI from 1987 to 2013, the FHFA expanded repeat-sales HPI from 1991 to 2013, the NAR median price index for single-family existing homes from 1968 to 2013, and the CoreLogic national median price index for all sales from 1976 to 2013. For inventory, the first row uses months of supply created by dividing homes vacant for sale from the Census Vacancy Survey by volume for single-family existing homes from the NAR. The second row just uses only the numerator, homes listed for sale from the Census Vacancy Survey, and adjusts the data for a linear time trend. The third row uses months of supply from 1982 to 2013 from the NAR. Robust standard errors are in parenthesis.

where c represents an MSA, $x_{c,t}$ is the vector of log months of supply, log price, and log sales volume in city c at time t, and φ_c is a city fixed effect. I use a Cholesky decomposition with the variables ordered so that months of supply is assumed to not depend contemporaneously on shocks to price or volume and price is assumed not to depend contemporaneously on shocks to volume. The results are robust as long as months of supply does not depend contemporaneously on volume.

To show the robustness of these result, Figure 31 estimates a two-lag vector autoregression model and a two-lag vector error correction (VEC) model on the national data sets used in the main text. The Cholesky ordering and VAR are the same as for the panel VAR except there is a single time series. The results look similar, although months of supply does not mean revert as quickly. Furthermore, the VEC model looks similar to the VAR model, which is reassuring if one is worried about the variables being cointegrated.

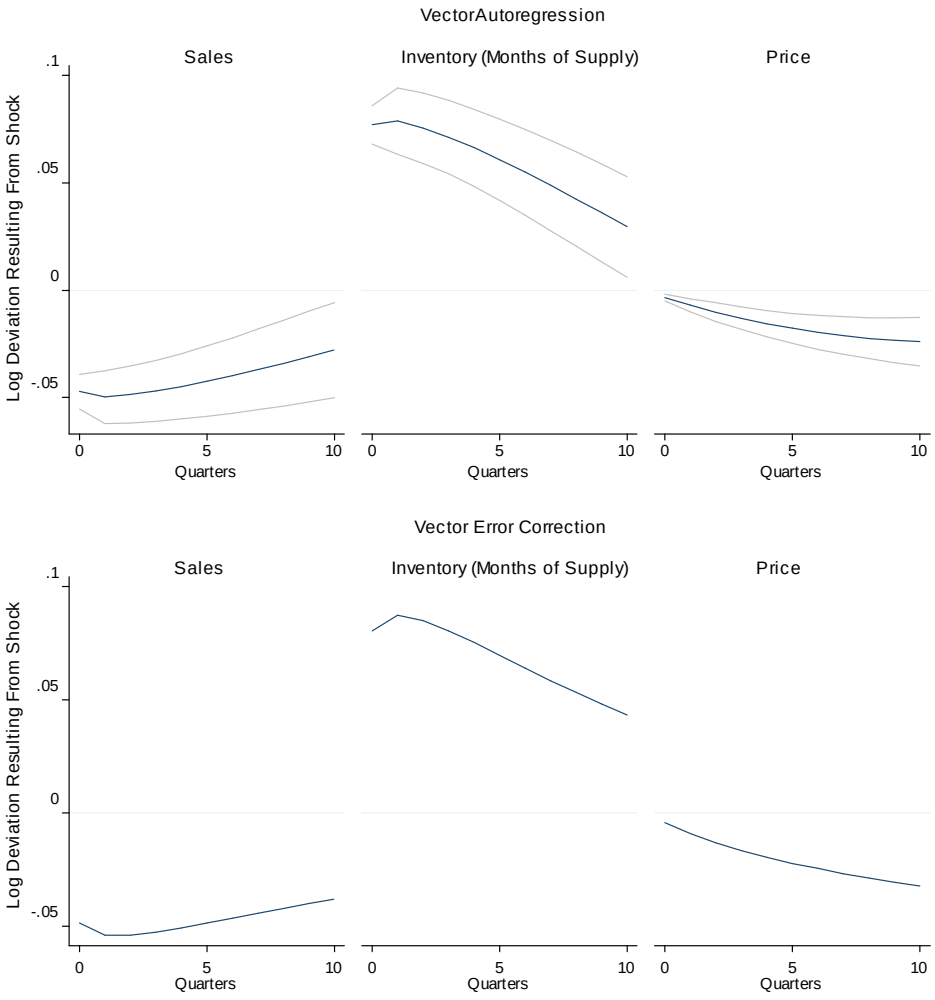

Figure 31: Impulse Response to Inventory Shock in VAR and VEC on National Data

Notes: The figures show orthogonalized impulse response functions to a months of supply shock from a two-lag vector autoregression model and vector error correction model of log months of supply, log price, and log sales volume. Price is the CoreLogic national HPI, sales is from the NAR single-family existing home sales series, and months of supply is from the Census Vacancy Survey and NAR, all from 1976-2013. All data are seasonally adjusted. The OIRFs are computed using a Cholesky decomposition with the variables ordered so that months of supply is assumed to not depend contemporaneously on shocks to price or volume and price is assumed to not depend contemporaneously on shocks to volume. The results are robust as long as months of supply is prior to volume in the Cholesky ordering. The blue line is the OIRF, and for the VAR the grey bands indicate 95% confidence intervals.

A.2.3 Buyer and Seller Entry

To show the robustness of Fact 4, which shows that entrants and sales move in opposite directions at peaks and troughs, Figure 32 shows the full data series calculated using the Census' homes for sale measure from 1968-2013 as well as the 2003 to 2013 period for months of supply from the National Association of Realtors. Although the patterns in previous cycles are not as dramatic,

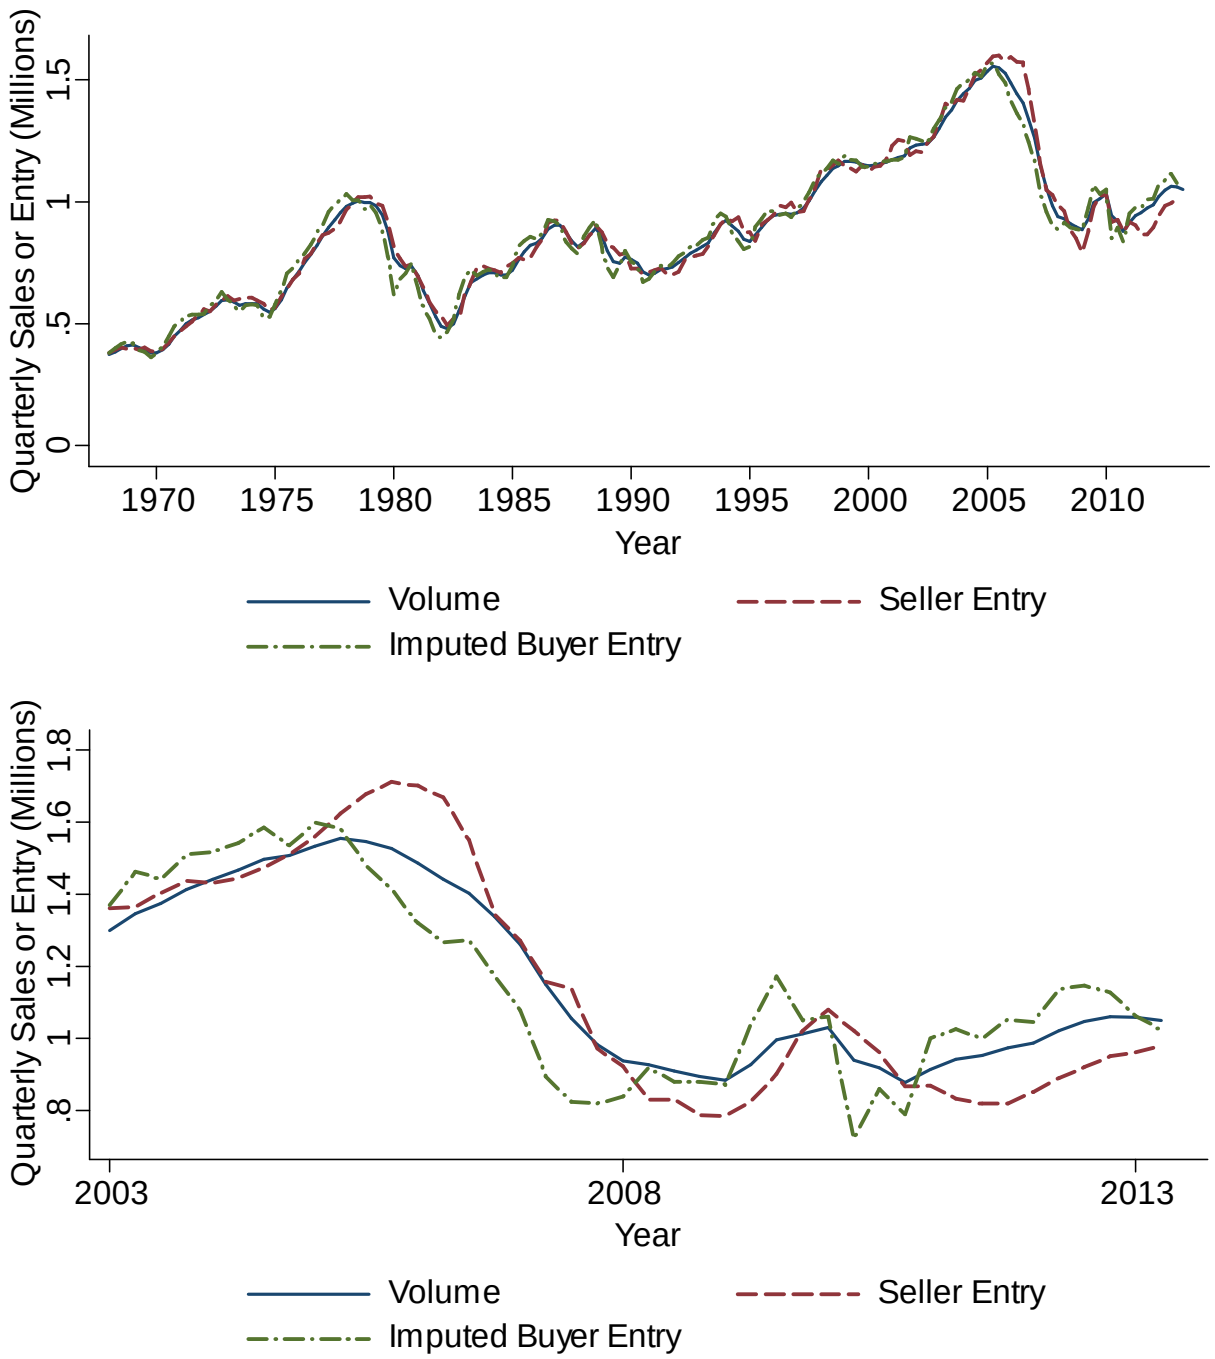

Figure 32: Sales, Entry: Full Series and NAR Data

Notes: Volume is raw data from the National Association of Realtors of sales of existing single-family homes at a seasonally-adjusted annual rate. The top panel shows seller entry using the stock of sellers measured by the Census Vacancy Survey, while the bottom panel shows the stock of sellers measured by the National Association of Realtors. Seller entry is computed as $\text{Entrants}_t = \text{Sellers}_t - \text{Sellers}_{t-1} + \text{Sales}_t$. Buyer entry is computed similarly, but since there is not a raw data series for the stock of buyers it is imputed using a simple Cobb-Douglas matching function $\frac{Sales}{S} = \xi \left(\frac{B}{S}\right)^{-.8}$ with the 0.8 elasticity from Genesove and Han (2012). In this figure, $\xi = 1$ in the top panel so that in a steady state there is 3 months of supply. Because the NAR series reports more homes listed for sale, $\xi = .5$ in the bottom panel to fit the same average market tightness as the Census series. All three series are smoothed using a three-quarter moving average in both figures.

Table 28: Cross-City Facts on Momentum, Inventory Volatility, and the Housing Phillips Curve

Dependent Var	SD of Annual Log Change in Months of Supply	Log Price Changes on Log Inventory Levels Regression Coefficient	Log Price Changes on Log Inventory Levels Regression R-Squared
Annual $\Delta \log(p)$	0.201*	-0.159**	0.920***
AR(1) Coefficient	(0.099)	(0.055)	(0.258)
N	42	42	42

Notes: * p < 0.05, ** p<0.01, *** p<0.001. Each column shows a regression of the indicated dependent variable calculated at the MSA level on the AR(1) coefficient from a regression of the annual change in log price on a one-year lag of itself, as in equation (1). These regressions thus show how the dependent variable varies across cities based on the degree of momentum the cities exhibit. The data used is the merged National Association of Realtors and CoreLogic data for 42 cities used in the panel VAR and described in Appendix A.1. Months of supply is directly from the NAR. The regression of the annual log change in price on log inventory levels is as in Equation (26). Robust standard errors are in parenthesis.

one can still see the same pattern. Furthermore, the pattern is clearly visible for the 2000s boom and bust in the NAR data.

Cross-City Facts Table 28 tests two of the model's predictions about the housing cycle facts in the cross-section of 42 cities used in the panel VAR analysis and described in Appendix A.1. Each column shows a regression in which the independent variable is a city-level measure of momentum: the AR(1) coefficient of the annual change in log price regressed on an annual lag of itself in quarterly CoreLogic data, as in equation (1). The first column shows that momentum co-varies positively with the standard deviation of the annual change in log months of supply, which comes from a separate data set from the National Association of Realtors. In the model, more momentum causes more re-timing of purchases and more inventory volatility. The second and third columns show the relationship between momentum and the regression coefficient and R-squared, respectively, from a regression of the annual change in log price on log months of supply at the midpoint of the year as in equation (26). Recall that in the data the coefficient is negative. Table 28 shows that the correlation between price changes and inventory levels is more negative and stronger the greater the degree of momentum in the city. This is consistent with the model, in which the housing Phillips curve arises due to the rapid adjustment of inventory and gradual adjustment of price in light of momentum.

A.3 Micro Evidence For Concave Demand

A.3.1 Binned Scatter Plots

Throughout the analysis I use binned scatter plots to visualize the structural relationship between list price relative to the reference list price and probability of sale. This section briefly describes how they are produced.

Recall that my econometric model is:

$$d_{h\ell t} = g\left(p_{h\ell t} - \tilde{p}_{h\ell t}\right) + \psi_{\ell t} + \varepsilon_{h\ell t} \tag{44}$$

where $p_{h\ell t} - \tilde{p}_{h\ell t}$ is equal to $f\left(z_{h\ell t}\right)$ in:

$$p_{h\ell t} = f\left(z_{h\ell t}\right) + \beta X_{h\ell t} + \xi_{\ell t} + u_{h\ell t}. \tag{45}$$

To create the IV binned scatter plots. I first estimate $f\left(z_{h\ell t}\right)$ by (45) and let $p_{h\ell t} - \tilde{p}_{h\ell t} = f\left(z_{h\ell t}\right)$. I drop the 1st and 99th percentiles of $p_{h\ell t} - \tilde{p}_{h\ell t}$ and ZIP-quarter cells with a single observation and create 25 indicator variables ζ_b corresponding to 25 bins q of $p_{h\ell t} - \tilde{p}_{h\ell t}$. I project sale within 13 weeks $d_{h\ell t}$ on fixed effects and the indicator variables:

$$d_{h\ell t} = \psi_{\ell t} + \zeta_b + \nu_{h\ell tq} \tag{46}$$

I visualize $g\left(\cdot\right)$ by plotting the average $p_{h\ell t} - \tilde{p}_{h\ell t}$ for each bin against the average $d_{h\ell t} - \psi_{\ell t}$ for each bin, which is equivalent to ζ_b.

A.3.2 Proof of Lemma 2

Recall that the Lemma assumes that:

$$\begin{aligned} & z_{h\ell t} \perp\!\!\!\perp \left(u_{h\ell t}, \varepsilon_{h\ell t}\right), \\ p_{h\ell t} \quad = \quad & f\left(z_{h\ell t}\right) + \zeta_{h\ell t} + \tilde{p}_{h\ell t}, \end{aligned}$$

$\zeta_{h\ell t} \perp\!\!\!\perp f(z_{h\ell t})$, and that the true regression function $g(\cdot)$ is a third-order polynomial. Because of the fixed effect $\xi_{h\ell t}$ in $\tilde{p}_{h\ell t}$, $\zeta_{h\ell t}$ can be normalized to be mean zero. Using the third-order polynomial assumption, the true regression function is:

$$\begin{aligned} g(p_{h\ell t} - \tilde{p}_{h\ell t}) &= E[d_{h\ell tq}|f(z_{h\ell t}) + \zeta_{h\ell t}, \psi_{\ell t}] = \beta_1 (f(z_{h\ell t}) + \zeta_{h\ell t}) \\ &\quad + \beta_2 (f(z_{h\ell t}) + \zeta_{h\ell t})^2 + \beta_3 (f(z_{h\ell t}) + \zeta_{h\ell t})^3 . \end{aligned}$$

However, $\zeta_{h\ell t}$ is unobserved, so I instead estimate:

$$\begin{aligned} E[d_{h\ell tq}|f(z_{h\ell t}), \psi_{\ell t}] &= \beta_1 f(z_{h\ell t}) + \beta_2 f(z_{h\ell t})^2 + \beta_3 f(z_{h\ell t})^3 \\ &\quad + \beta_1 E[\zeta_{h\ell t}|f(z_{h\ell t})] + 2\beta_2 E[f(z_{h\ell t})\zeta_{h\ell t}] + \beta_2 E\left[\zeta_{h\ell t}^2|f\right] \\ &\quad + 3\beta_3 f(z_{h\ell t}) E\left[\zeta_{h\ell t}^2|f\right] + 3\beta_3^2 f(z_{h\ell t}) E[\zeta_{h\ell t}|f] + \beta_3 E\left[\zeta_{h\ell t}^3|f\right] . \end{aligned}$$

However, because $\zeta_{h\ell t} \perp\!\!\!\perp f(z_{h\ell t})$, $E[\zeta_{h\ell t}|f(z_{h\ell t})] = 0$, $E[f(z_{h\ell t})\zeta_{h\ell t}] = 0$, and $E\left[\zeta_{h\ell t}^2|f\right]$ and $E\left[\zeta_{h\ell t}^3|f\right]$ are constants. The $\beta_2 E\left[\zeta_{h\ell t}^2|f\right]$ and $\beta_3 E\left[\zeta_{h\ell t}^3|f\right]$ terms will be absorbed by the fixed effects $\psi_{\ell t}$, leaving:

$$E[d_{h\ell tq}|f(z_{h\ell t}), \psi_{\ell t}] = \beta_1 f(z_{h\ell t}) + \beta_2 f(z_{h\ell t})^2 + \beta_3 f(z_{h\ell t})^3 + 3\beta_3 f(z_{h\ell t}) E\left[\zeta_{h\ell t}^2|f\right]$$

Thus when one estimates $g(\cdot)$ by a cubic polynomial of $f(z_{h\ell t})$,

$$d_{h\ell tq} = \gamma_1 f(z_{h\ell t}) + \gamma_2 f(z_{h\ell t})^2 + \gamma_3 f(z_{h\ell t})^3 + \psi_{\ell t} + \varepsilon_{h\ell t},$$

one recovers $\gamma_1 = \beta_1 + 3\beta_3 E\left[\zeta_{h\ell t}^2|f\right]$, $\gamma_2 = \beta_2$, and $\gamma_3 = \beta_3$, so the true second- and third-order terms are recovered.

A.3.3 Instrumental Variable Robustness and Specification Tests

This section provides robustness and specification tests for the IV estimates described in Section 1.3.

Figure 33 shows the reduced-form relationship between the instrument and outcome variable

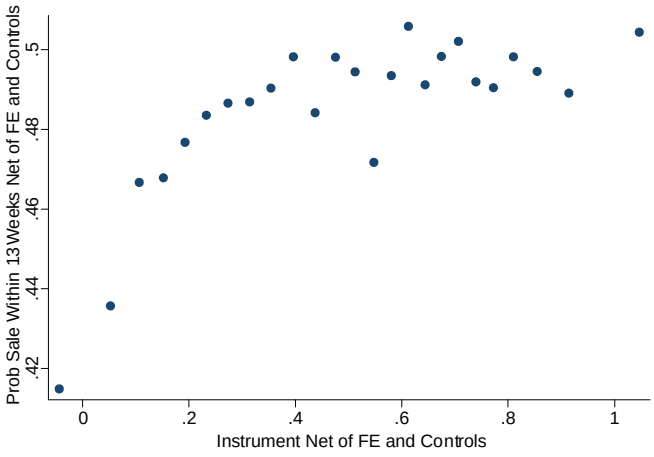

Figure 33: Reduced-Form Relationship Between the Instrument and the Outcome Variable

Notes: This figure shows the reduced-form relationship between the instrument on the x-axis and the probability of sale within 13 weeks on the y axis. Both are residualized against quarter x ZIP fixed effects and the repeat-sales and hedonic predicted prices and the means are added back in. This is the basic concave relationship that the IV approach uses, although the downward-sloping first stage flips the x-axis.

when both are residualized against fixed effects and the repeat-sales and hedonic predicted price. The estimates presented in the main text rescale the instrument axis into price (and in the process flip the x axis), but the basic concave relationship between probability of sale and appreciation since purchase is visible in the reduced form.

Figure 34 shows IV binned scatter plots when the y-axis is rescaled to a logarithmic scale so that the slope represents the elasticity of demand, which does not alter the finding of concavity.

Figure 35 shows third-order polynomial fits varying the number of weeks that a listing needs to sell within to count as a sale from six weeks to 26 weeks. Concavity is evident regardless of the deadline used.

Figure 36 shows the IV binned scatter plot and a third-order polynomial fit when the sample is limited to transactions and the prices are measured in list prices rather than transaction prices. Substantial concavity is still present, assuaging concerns that the concavity in list prices may not translate into a strategic complementarity in transaction prices. The upward slope in the middle of the figure is not statistically significant.

Tables 29, 30, 31, and 32 present various robustness and specification tests of the main IV specification in Panel B of Table 3. Each row in the tables represents a separate regression, with the specifications described in the main text. Coefficients for a three-segment spline in the log relative

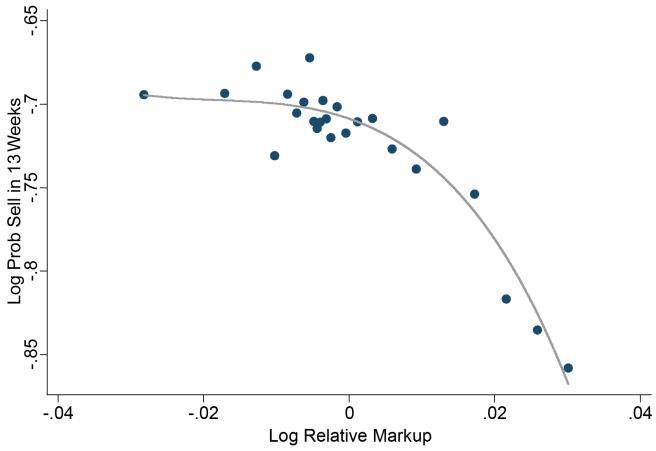

Figure 34: Instrumental Variable Estimates With Probability of Sale Axis in Logs

Notes: The figure shows a binned scatter plot of the log of probability of sale within 13 weeks net of fixed effects (with the average probability of sale within 13 weeks added in) against the estimated log relative markup. It also shows an overlaid cubic fit of the relationship, as in equation (3). To create the figure, a first stage regression of the log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and repeat sales and hedonic log predicted prices, as in (6), is estimated by OLS. The predicted value of the polynomial of the instrument is used as the relative markup. The figure splits the data into 25 equally-sized bins of this estimated relative markup and plots the mean of the estimated relative markup against the mean of the probability of sale within 13 weeks net of fixed effects for each bin. The y-axis is rescaled into logs after means are calculated and the cubic fit is estimated because the outcome variable is binary. Before binning, the 1st and 99th percentiles of the log sale price residual and any observations fully absorbed by fixed effects are dropped. The entire procedure is weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. The sample is limited to the IV subsample of homes that are not sales of foreclosures or short sales, sales of homes with negative appreciation since the seller purchased, or sales by investors who previously purchased with all cash. N = 111,293 observations prior to dropping the 1st and 99th percentiles and unique zip-quarter cells.

markup, the difference between the highest and lowest tercile coefficients, and a bootstrapped 95 percent confidence interval for the difference are reported. In some specifications the bootstrapped confidence intervals widen when the sample size is reduced to the point that the results are no longer significant, and the middle tercile slope can be sensitive because the middle third of the data tends to correspond to a very small range of log relative markups and is thus noisily estimated. Nonetheless, the robustness and specification checks show evidence of significant concavity.

Table 29 evaluates the exclusion restriction that unobserved quality is independent of when a seller purchased. The first two specifications add a linear trend in date of purchase or time since purchase in $X_{h\ell t}$ along with the two predicted prices, thus accounting for any variation in unobserved quality that varies linearly in date of purchase or time since purchase. The next three rows limit the sample to homes purchased before the bust (before 2005), after 1994, and in a window

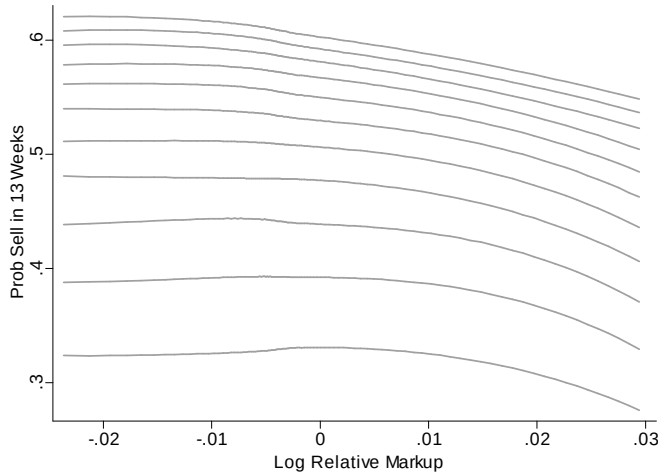

Figure 35: Instrumental Variable Estimates: Varying The Sell-By Date

Notes: The figure shows third-order polynomial fits of equation (3) for the probability of sale by eleven different deadlines (6, 8, 10, 12, 14, 16, 18, 20, 22, 24, and 26 weeks) net of fixed effects (with the average probability of sale added in) against the estimated log relative markup. To create the figure, a first stage regression of the log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and repeat sales and hedonic log predicted prices, as in (6), is estimated by OLS. The predicted value of the polynomial of the instrument is used as the relative markup before equation (3) is run. The entire procedure is weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. The sample is limited to the IV subsample of homes that are not sales of foreclosures or short sales, sales of homes with negative appreciation since the seller purchased, or sales by investors who previously purchased with all cash. N = 111,293 observations prior to dropping the 1st and 99th percentiles and unique zip-quarter cells.

from 1995 to 2004. Finally, the last two rows add linear time trends to the purchased before 2005 sample. In all cases, the bootstrapped 95 percent confidence intervals continue to show significant concavity.

Table 30 shows various specification checks. The first set of regressions limit the analysis to ZIP-quarter cells with at least 15 and 20 observations to evaluate whether small sample bias in the estimated fixed effect $\xi_{h\ell t}$ could be affecting the results. In both cases, the results appear similar to the full sample and the bootstrapped confidence interval shows a significant difference between the highest and lowest terciles, which suggests that bias in the estimation of the fixed effects is not driving the results. The second set introduces $X_{h\ell t}$, the vector of house characteristics that includes the repeat-sales and hedonic predicted prices, as a quadratic and cubic function instead of linearly. It does not appear that assumed linearity of these characteristics is driving the results. Finally, the third set considers different specifications for the flexible function of the instrument $f(z_{h\ell t})$ in the first stage. Again, changing the order of $f(\cdot)$ does not appear to have a significant

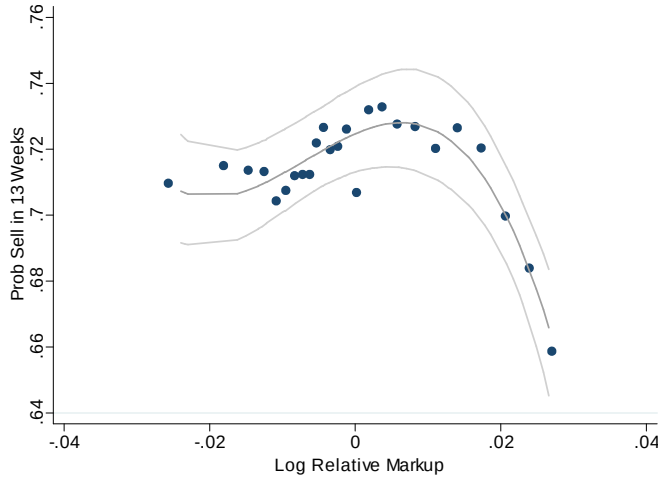

Figure 36: Instrumental Variable Estimates: Transaction Prices

Notes: The figure shows a binned scatter plot of the probability of sale within 13 weeks net of fixed effects (with the average probability of sale within 13 weeks added in) against the estimated log relative markup measured using transaction prices rather than list prices. It also shows an overlaid cubic fit of the relationship, as in equation (3). To create the figure, a first stage regression of the log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and repeat sales and hedonic log predicted prices, as in (6), is estimated by OLS. The predicted value of the polynomial of the instrument is used as the relative markup. The figure splits the data into 25 equally-sized bins of this estimated relative markup and plots the mean of the estimated relative markup against the mean of the probability of sale within 13 weeks net of fixed effects for each bin. Before binning, the 1st and 99th percentiles of the log sale price residual and any observations fully absorbed by fixed effects are dropped. The entire procedure is weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. The sample is limited to the IV subsample of homes that are not sales of foreclosures or short sales, sales of homes with negative appreciation since the seller purchased, or sales by investors who previously purchased with all cash. To obtain transaction prices, the sample is also limited to homes that transact. The grey bands indicate a pointwise 95-percent confidence interval for the cubic fit created by block bootstrapping the entire procedure on 35 ZIP-3 clusters. N = 74,299 observations prior to dropping the 1st and 99th percentiles and unique zip-quarter cells.

effect on concavity.

Table 31 shows various robustness checks. These include:

- House Characteristic Controls: Includes a third-order polynomial in age, log square feet, bedrooms, and bathrooms in $X_{h\ell t}$.
- Alternate Time To Sale Definition: Instead of measuring time to sale as first listing to the filing of the deed transfer request, this specification measures time to sale as first listing to the first of the deed transfer request or the last listing.
- 18 and 10 Weeks to Sale: Tests the robustness to the horizon for having sold.

Table 29: IV Robustness 1: Controls for Time Since Purchase

Dependent Var: Sell Within 13 Weeks						
Specification	Tercile Spline Coefficients			Difference		
(Details In Text)	Lowest	Middle	Highest	High - Low	Bootstrapped 95% CI	Obs
Linear Trend in	-0.391	1.009	-2.826***	-2.435**	[-5.014, -1.994]	111,293
Date of Purchase	(0.403)	(4.134)	(0.810)	(0.836)		
Linear Trend in	-0.396	1.109	-2.849***	-2.453**	[-5.065, -1.999]	111,293
Time Since Purchase	(0.407)	(4.240)	(0.824)	(0.848)		
Purchased Pre 2005	-0.038	-0.357	-4.11	-4.072	[-13.155, -2.282]	102,642
	(0.374)	(1.907)	(4.115)	(4.116)		
Purchased Post 1994	-0.702	-0.095	-2.503***	-1.800	[-4.375, -1.790]	93,080
	(0.769)	(0.952)	(0.676)	(1.025)		
Purchased 1995-2004	-0.407	-0.816	-4.252	-3.845	[-15.221, -2.432]	84,429
	(0.782)	(2.193)	(5.149)	(5.443)		
Pre 2005 With Trend	-0.042	-0.152	-3.183	-3.140	[-7.326, -1.871]	102,642
in Date of Purchase	(0.306)	(0.999)	(2.157)	(2.124)		
Pre 2005 With Trend	-0.042	-0.156	-3.21	-3.168	[-7.624, -1.893]	102,642
in Time Since Purchase	(0.307)	(1.016)	(2.238)	(2.204)		

Notes: * p < 0.05, ** p<0.01, *** p<0.001. Each row shows regression coefficients when g(.) in equation (3) is approximated using a three-segment linear spline with an equal fraction of the data in each segment. A first stage regression of log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and log predicted price using both a repeat-sales and a hedonic methodology, as in (6), is estimated by OLS. The predicted value of the polynomial of the instrument is computed is used as the relative markup in equation (3), which is estimated by OLS. The sample is restricted to non-REOs, non-short sales, properties with positive appreciation since purchase, and properties not previously purchased with all cash (investors). The entire procedure weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. Before creating the spline, the 99th and 1st percentiles of the relative markup are dropped, as are any observations fully absorbed by fixed effects. In addition to the regression coefficients, the difference between the highest and lowest tercile of the spline is reported. Standard errors and the 95 percent confidence interval for the difference between the first and third terciles are computed by block bootstrapping the entire procedure on 35 ZIP-3 clusters. The number of observations listed is prior to dropping observations that are unique to a ZIP-quarter cell and the 1st and 99th percentiles. The appendix text details each specification.

- Non-REO House Price Index for Predicted Price: Uses house price indices for the predicted price that does not include REOs.
- Nearby REOs: Controls for third order polynomial in number of REOs within 1/4 mile from 2006 to 2013 to account for quality correlated with distressed sales at the sup ZIP code level. The results are similar if one varies the radius considered from .1 to 1 mile or counts REOs within a year of the listing rather than from 2006 to 2013.
- No Weights: Does not weight observations by the inverse standard deviation of the repeat-sales house price index prediction error at the ZIP level.

Table 30: IV Robustness 2: Specification Checks

Dependent Var: Sell Within 13 Weeks						
Specification	Tercile Spline Coefficients			Difference		
(Details In Text)	Lowest	Middle	Highest	High - Low	Bootstrapped 95% CI	Obs
Only FE Cells With	-0.068	-0.295	-2.575***	-2.507***	[-4.31, -1.844]	64,236
At Least 15 Obs	(0.455)	(2.731)	(0.705)	(0.667)		
Only FE Cells With	-0.285	-2.024	-3.163**	-2.878	[-6.835, -1.265]	44,795
At Least 20 Obs	(1.052)	(4.378)	(1.192)	(1.480)		
Predicted Prices	-0.225	-0.425	-2.248***	-2.023***	[-3.324, -1.307]	111,293
Introduced as Quadratic	(0.353)	(2.154)	(0.503)	(0.534)		
Predicted Prices	-0.228	-0.361	-2.314***	-2.086***	[-3.313, -1.344]	111,293
Introduced as Cubic	(0.355)	(2.014)	(0.500)	(0.532)		
Linear Fn	-0.440	0.260	-4.511***	-4.071***	[-6.632, -2.926]	111,293
of Instrument	(0.503)	(0.603)	(0.982)	(0.900)		
Quadratic Fn	-0.466	0.262	-4.301**	-3.835**	[-7.602, -1.998]	111,293
of Instrument	(0.594)	(0.601)	(1.438)	(1.466)		
Quartic Fn	-0.279	0.069	-2.171***	-1.892***	[-3.237, -1.235]	111,293
of Instrument	(0.302)	(3.092)	(0.533)	(0.536)		
Quintic Fn	-0.235	-0.560	-2.120***	-1.885***	[-3.340, -1.150]	111,293
of Instrument	(0.404)	(2.248)	(0.518)	(0.573)		

Notes: * p < 0.05, ** p<0.01, *** p<0.001. Each row shows regression coefficients when g(.) in equation (3) is approximated using a three-segment linear spline with an equal fraction of the data in each segment. A first stage regression of log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and log predicted price using both a repeat-sales and a hedonic methodology, as in (6), is estimated by OLS. The predicted value of the polynomial of the instrument is computed is used as the relative markup in equation (3), which is estimated by OLS. The sample is restricted to non-REOs, non-short sales, properties with positive appreciation since purchase, and properties not previously purchased with all cash (investors). The entire procedure weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. Before creating the spline, the 99th and 1st percentiles of the relative markup are dropped, as are any observations fully absorbed by fixed effects. In addition to the regression coefficients, the difference between the highest and lowest tercile of the spline is reported. Standard errors and the 95 percent confidence interval for the difference between the first and third terciles are computed by block bootstrapping the entire procedure on 35 ZIP-3 clusters. The number of observations listed is prior to dropping observations that are unique to a ZIP-quarter cell and the 1st and 99th percentiles. The Appendix text details each specification.

- No Possibly Problematic Observations: A small number of listings are matched to multiple property IDs and I use an algorithm described in Appendix A.1 to guess of which is the relevant property ID. Additionally, there are spikes in the number of listings in the Altos data for a few dates, which I have largely eliminated by dropping listings that do not match to a DataQuick property ID. Despite the fact that these two issues affect a very small number of observations, this specification drops both types of potentially problematic observations to

Table 31: IV Robustness 3: Other Robustness Tests

Dependent Var: Sell Within 13 Weeks Unless Otherwise Indicated						
Specification	Tercile Spline Coefficients			Difference		
(Details In Text)	Lowest	Middle	Highest	High - Low	Bootstrapped 95% CI	Obs
House Characteristic	-0.366	0.350	-2.804***	-2.438***	[-3.938, -1.567]	107,176
Controls	(0.359)	(1.953)	(0.589)	(0.593)		
Alternate Time	-0.587	0.186	-2.104***	-1.516**	[-2.772, -0.660]	111,293
to Sale Defn	(0.357)	(1.772)	(0.506)	(0.561)		
Dep Var: 18 Weeks	-0.240	-0.974	-2.140***	-1.900***	[-3.124, -1.159]	111,293
	(0.339)	(1.685)	(0.524)	(0.537)		
Dep Var: 10 Weeks	-0.049	0.611	-2.302***	-2.252***	[-3.742, -1.359]	111,293
	(0.413)	(1.842)	(0.549)	(0.588)		
No REO HPI	-0.139	0.391	-1.248***	-1.110*	[-2.231, -0.210]	108,081
For Predicted Price	(0.505)	(2.549)	(0.253)	(0.541)		
Nearby REOs	-0.393	1.012	-2.827***	-2.434***	[-4.381,-1.332]	111,293
	(0.402)	(4.136)	(0.810)	(0.836)		
No Weights	-0.153	1.109	-2.046***	-1.892***	[-3.015, -1.062]	111,293
	(0.399)	(1.394)	(0.511)	(0.505)		
No Poss Problematic Obs	-0.280	0.246	-2.355***	-2.075***	[-3.526, -1.47]	107,865
	(0.325)	(1.768)	(0.588)	(0.551)		
First Listed 2008-6/2010	0.273	-0.871	-2.967***	-3.240***	[-5.684, -2.174]	57,384
	(0.507)	(4.533)	(0.766)	(0.902)		
First Listed 7/2010-2013	-1.034	1.265	-2.039*	-1.005	[-3.030, 1.199]	53,909
	(0.907)	(1.429)	(0.894)	(1.150)		
Bay Area	-0.324	1.292	-4.430	-4.105	[-10.067, -1.505]	29,108
	(1.348)	(2.284)	(2.786)	(2.892)		
Los Angeles	0.152	0.251	-2.172**	-2.325**	[-4.203, -1.144]	68,749
	(0.656)	(2.542)	(0.688)	(0.813)		
San Diego	-3.344	1.068	-2.683	0.661	[-54.876, 2.189]	13,436
	(11.060)	(2.741)	(9.616)	(20.092)		

Notes: * p < 0.05, ** p<0.01, *** p<0.001. Each row shows regression coefficients when g(.) in equation (3) is approximated using a three-segment linear spline with an equal fraction of the data in each segment. A first stage regression of log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and log predicted price using both a repeat-sales and a hedonic methodology, as in (6), is estimated by OLS. The predicted value of the polynomial of the instrument is computed is used as the relative markup in equation (3), which is estimated by OLS. The sample is restricted to non-REOs, non-short sales, properties with positive appreciation since purchase, and properties not previously purchased with all cash (investors). The entire procedure weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. Before creating the spline, the 99th and 1st percentiles of the relative markup are dropped, as are any observations fully absorbed by fixed effects. In addition to the regression coefficients, the difference between the highest and lowest tercile of the spline is reported. Standard errors and the 95 percent confidence interval for the difference between the first and third terciles are computed by block bootstrapping the entire procedure on 35 ZIP-3 clusters. The number of observations listed is prior to dropping observations that are unique to a ZIP-quarter cell and the 1st and 99th percentiles. The Appendix text details each specification.

Table 32: IV Robustness 4: Transactions Only

Dependent Variable	Tercile Spline Coefficients			Difference		
(Details In Text)	Lowest	Middle	Highest	High - Low	Bootstrapped 95% CI	Obs
Sell Within 13 Weeks	0.254	4.376	-2.169***	-2.423**	[-4.243, -0.825]	74,299
	-0.778	-2.500	-0.486	-0.923		
Weeks on Market	(25.263)	-102.910*	62.734***	87.996***	[58.917, 143.191]	74,299
	-20.157	-50.969	-11.02	-19.655		
Weeks on Market	(18.203)	-126.848*	56.316***	74.519***	[51.993, 110.053]	74,299
Alternate Defn	-16.871	-58.846	-10.094	-15.106		

Notes: * p < 0.05, ** p<0.01, *** p<0.001 .Each row shows regression coefficients when g(.) in equation (3) is approximated using a three-segment linear spline with an equal fraction of the data in each segment. A first stage regression of log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and log predicted price using both a repeat-sales and a hedonic methodology, as in (6), is estimated by OLS. The predicted value of the polynomial of the instrument is computed is used as the relative markup in equation (3), which is estimated by OLS. The sample is restricted to non-REOs, non-short sales, properties with positive appreciation since purchase, and properties not previously purchased with all cash (investors). The entire procedure weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. Before creating the spline, the 99th and 1st percentiles of the relative markup are dropped, as are any observations fully absorbed by fixed effects. In addition to the regression coefficients, the difference between the highest and lowest tercile of the spline is reported. Standard errors and the 95 percent confidence interval for the difference between the first and third terciles are computed by block bootstrapping the entire procedure on 35 ZIP-3 clusters. The number of observations listed is prior to dropping observations that are unique to a ZIP-quarter cell and the 1st and 99th percentiles. The Appendix text details each specification.

show that they do not affect results.

- By Time Period: I split the data into two time periods, February 2008 to June 2010 and July 2010 to February 2013.

- By MSA: Separate regressions for the San Francisco Bay, Los Angeles, and San Diego areas.

The results continue to show concavity, although in some specifications it is weakened by the smaller sample size and no longer significant. In particular, in San Diego the confidence intervals are so wide that nothing can be inferred. Additionally, both when predicted prices are computed using non-REO house price indices (which are much noisier) and in the second half of the sample, the result is slightly weakened and no longer significant.

Finally, Table 32 shows results for the subset of homes that transact for three different outcome variables. First, it shows the main sale within 13 weeks outcome, for which the concavity is still significant. The second two specifications show results using weeks on the market as the outcome variable, and so concavity is indicated by a positive difference between the highest and lowest

terciles. For both the baseline and alternate weeks on the market definitions, there is significant concavity in the IV specifications that indicates that increasing the list price by one percent increases time on the market by 1.1 to 1.5 weeks.

A.3.4 Ordinary Least Squares

An alternative to IV is to assume that there is no unobserved quality and thus no need for an instrument. This ordinary least squares approach implies that:

$$\tilde{p}_{h\ell t} = \xi_{\ell t} + \beta X_{h\ell t}$$

and so $p_{h\ell t} - \tilde{p}_{h\ell t}$ is equal to the regression residual $\eta_{h\ell t}$ in:

$$p_{h\ell t} = \xi_{\ell t} + \beta X_{h\ell t} + \eta_{h\ell t}, \tag{47}$$

which can be estimated in a first stage and plugged into the second stage equation:

$$d_{h\ell t} = g\left(\eta_{h\ell t}\right) + \psi_{\ell t} + \varepsilon_{h\ell t}.$$

Given the importance of unobserved quality, this is likely to provide significantly biased results, but it is worth considering as a benchmark as discussed in the main text. This section provides additional OLS results to show that the findings in Panel A of Table 3 are robust.

Because the OLS sample may include distressed sales, I take a conservative approach and include fixed effects at the ZIP × quarter × distress status level. Distressed status is defined as either non-distressed, REO, or a short sale (or withdrawn listing subsequently foreclosed upon). The results would look similar if ZIP × quarter fixed effects were used and an additive categorical control for distressed status were included in $X_{h\ell t}$.

First, Figure 37 shows binned scatter plots for OLS for all listings, transactions only, and the IV subsample. In each, a clear pattern of concavity is visible, but as discussed in the main text, the upward slope on the left indicates the presence of substantial unobserved quality—particularly among homes that do not sell—and thus the need for an instrument.[123]

[123] An alternative explanation is that in the later years of my sample I do not have follow-up data on foreclosures, so

Table 33: Ordinary Least Squares Robustness

Dependent Var: Weeks on Market						
Specification	Tercile Spline Coefficients			Difference		
(Details In Text)	Lowest	Middle	Highest	High - Low	Bootstrapped 95% CI	Obs
House Characteristic	0.155***	-0.293***	-0.553***	-0.708***	[-0.835, -0.616]	414,525
Controls	(0.034)	(0.047)	(0.034)	(0.058)		
Alternate Time	0.143***	-0.529***	-0.496***	-0.639***	[-0.762, -0.533]	431,830
to Sale Defn	(0.030)	(0.093)	(0.042)	(0.060)		
Dep Var: 18 Weeks	0.195***	-0.419***	-0.496***	-0.692***	[-0.802, -0.59]	431,830
	(0.031)	(0.090)	(0.034)	(0.056)		
Dep Var: 10 Weeks	0.122***	-0.536***	-0.453***	-0.575***	[-0.692, -0.473]	431,830
	(0.030)	(0.088)	(0.039)	(0.056)		
Hedonic Predicted	0.128***	-0.255***	-0.458***	-0.587***	[-0.724, -0.427]	665,560
Price Only	(0.038)	(0.049)	(0.043)	(0.076)		
Low REO ZIPs	0.152**	-0.943***	-0.465***	-0.618***	[-0.847, -0.513]	180,517
	(0.052)	(0.086)	(0.050)	(0.095)		
Low Short Sale ZIPs	0.060	-1.063***	-0.405***	-0.465**	[-0.773, -0.272]	113,196
	(0.073)	(0.124)	(0.084)	(0.147)		
No REO or Short Sale	0.322***	-0.738***	-0.525***	-0.846***	[-1.052, -0.684]	221,013
	(0.043)	(0.175)	(0.068)	(0.096)		
Transactions Only	0.012	-0.419***	-0.510***	-0.522***	[-0.607, -0.46]	318,842
	(0.029)	(0.045)	(0.021)	(0.038)		
IV Subsample	0.322***	-1.110***	-0.533***	-0.854***	[-1.141, -0.65]	111,293
	(0.073)	(0.154)	(0.073)	(0.137)		
Only FE Cells With	0.156***	-0.610***	-0.496***	-0.652***	[-0.819, -0.563]	273,100
At Least 20 Obs	(0.034)	(0.117)	(0.057)	(0.072)		
Predicted Prices	0.144***	-0.405***	-0.534***	-0.679***	[-0.777, -0.588]	431,830
Introduced as Cubic	(0.033)	(0.039)	(0.026)	(0.050)		
First Listed 2008-7/2010	0.046	-0.616***	-0.522***	-0.567***	[-0.685, -0.467]	233,115
	(0.032)	(0.086)	(0.041)	(0.057)		
First Listed 7/2010-2013	0.317***	-0.375***	-0.415***	-0.731***	[-0.873, -0.642]	198,715
	(0.041)	(0.099)	(0.034)	(0.060)		
Bay Area	0.170**	-0.421***	-0.565***	-0.735***	[-0.879, -0.571]	115,223
	(0.065)	(0.097)	(0.058)	(0.081)		
Los Angeles	0.186***	-0.555***	-0.440***	-0.626***	[-0.79, -0.479]	256,816
	(0.048)	(0.109)	(0.044)	(0.081)		
San Diego	0.133***	-0.275*	-0.578***	-0.710***	[-0.748, -0.65]	59,791
	(0.013)	(0.134)	(0.039)	(0.027)		

Notes: * p < 0.05, ** p<0.01, *** p<0.001. Each row shows regression coefficients when g(.) in equation (3) is approximated using a three-segment linear spline with an equal fraction of the data in each segment. OLS estimates the relative markup based on a first stage, equation (47), and plugs in the estimate relative markup into equation (3). The fixed effects at the quarter of initial listing x ZIP x distress status level. Distress status corresponds to three groups: normal sales, REOs (sales of foreclosed homes and foreclosure auctions), and short sales (cases where the transaction was less than the amount outstanding on the loan and withdrawals that are subsequently foreclosed on in the next two years). Both procedures are weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. Before creating the spline, the 99th and 1st percentiles of the relative markup are dropped, as are any observations fully absorbed by fixed effects. In addition to the regression coefficients, the difference between the highest and lowest tercile of the spline is reported. Standard errors and the 95 percent confidence interval for the difference between the first and third terciles are computed by block bootstrapping the entire procedure on 35 ZIP-3 clusters. The number of observations listed is prior to dropping observations that are unique to a ZIP-quarter cell and the 1st and 99th percentiles. The Appendix text details each specification.

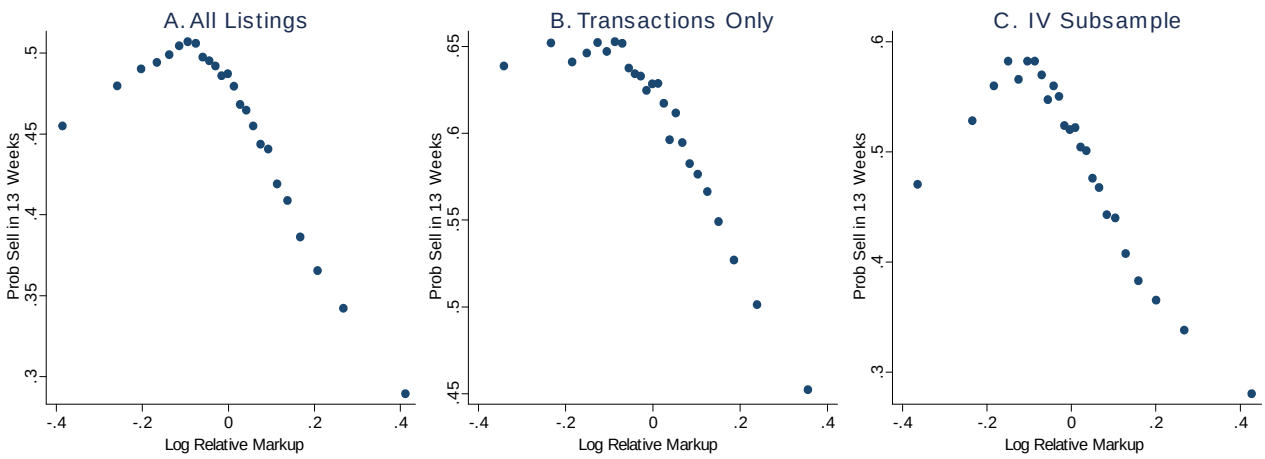

Figure 37: The Effect of List Price on Probability of Sale: Ordinary Least Squares

Notes: Each panel shows a binned scatter plot of the probability of sale within 13 weeks against the log relative markup. OLS assumes no unobserved quality. To create the figure, a first stage regression of log list price on fixed effects at the ZIP x first quarter of listing level x seller distress status level and repeat sales and hedonic log predicted prices, as in (47), is estimated by OLS. The residual is used as the relative markup in equation (3), which is estimated by OLS. The figure splits the data into 25 equally-sized bins of the estimated relative markup and plots the mean of the estimated relative markup against the log of the mean of the probability of sale within 13 weeks net of fixed effects for each bin. Before binning, the 1st and 99th percentiles of the log sale price residual and any observations fully absorbed by fixed effects are dropped. The entire procedure is weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. Observation counts prior to dropping the 1st and 99th percentiles are 431,830 for panel A (all listings with a prior observed sale), 318,832 for panel B (listings with a prior observed sale that lead to transactions), and 111,293 for panel C (IV sample).

Table 33 shows a number of robustness and specification checks. Those different from the IV specification checks described previously are:

- House Characteristic Controls: As with IV, this includes a third-order polynomial in age, log square feet, bedrooms, and bathrooms, but it also includes additive fixed effects for quintiles of the time since purchase distribution in $X_{h\ell t}$.
- Hedonic predicted price only: Drops the repeat-sales house price index from $X_{h\ell t}$. This expands the sample to all listings in the data rather than only those with a prior observed sale.
- Low REO ZIPs: Only includes ZIP codes with less than 20 percent REO sale shares from 2008 to 2013. (REO is a sale of a foreclosed property.)

some withdrawn short sales are counted as non-distressed. This may explain some of the upward slope, as the upward slope is concentrated in non-withdrawn properties, high short sale ZIP codes, and the later years of my sample.

- Low Short ZIPs: Only includes ZIP codes with less than 20 percent short sale shares from 2008 to 2013. (A short sale occurs when a homeowner sells their house for less than their outstanding mortgage balance and must negotiate the sale with their lender.)
- No REO or Short Sale: Drops REOs, short sales, and withdrawn sales subsequently foreclosed upon homes, thus only leaving non-distressed sales.
- Transactions only: Drops houses withdrawn from the market.
- IV Subsample: Drops homes with negative appreciation since purchase, REOs, and homes previously purchased with all cash.

All specifications show significant concavity.

A.3.5 Monte Carlo Assessment of Bias From Other Sources of Markup Variation

This appendix presents Monte Carlo simulations to assess the degree of bias from other sources of variation in the relative markup entering $g(\cdot)$ nonlinearly.

To do so, for each house in the IV sample I simulate $d_{h\ell t}$ using an assumed $g(\cdot)$ and an estimated $\psi_{h\ell t}$:

$$d_{h\ell t} = g\left(p_{h\ell t} - \tilde{p}_{h\ell t}\right) + \psi_{\ell t} + \varepsilon_{h\ell t}.$$

However, rather than assuming $p_{h\ell t} - \tilde{p}_{h\ell t} = f\left(z_{h\ell t}\right)$, I let $p_{h\ell t} - \tilde{p}_{h\ell t} = f\left(z_{h\ell t}\right) + \zeta_{h\ell t}$ and report results for different parameterizations for the other sources of relative markup variation $\zeta_{h\ell t}$.

Specifically, I follow a five step procedure 1,000 times and report the average values:

1. Based on first stage, calculate $p_{h\ell t} - \tilde{p}_{h\ell t} = f\left(z_{h\ell t}\right)$.
2. Estimate $\psi_{\ell t}$ given $g(\cdot)$
3. Draw $\zeta_{h\ell t}$. Using the known $g(\cdot)$, calculate $g\left(f\left(z_{h\ell t}\right) + \zeta_{h\ell t}\right) + \psi_{\ell t}$
4. $d_{h\ell t}$ is Bernoulli: a house sells with probability $g\left(f\left(z_{h\ell t}\right) + \zeta_{h\ell t}\right) + \psi_{\ell t}$
5. Run the estimator of interest on the simulated data.

Table 34: Monte Carlo Simulations: Other Sources of Markup Variation Independent of Instrument

3-Part Spline	SD of $\zeta_{h\ell t}$			
Coef Estimates	0.000	0.01	0.02	0.04
Lowest Tercile	-0.215	-0.851	-1.110	-2.486
	(0.371)	(0.571)	(0.422)	(0.379)
Middle Tercile	0.201	-0.768	-1.145	-3.002
	(1.037)	(1.249)	(1.090)	(1.008)
Highest Terclie	-2.219	-1.861	-2.077	-3.386
	(0.264)	(0.605)	(0.816)	(0.267)
High - Low	-2.006	-1.010	-0.968	-0.900
	(0.429)	(1.068)	(1.032)	(0.443)

Notes: Each column shows the mean and standard deviation over 1,000 Monte Carlo simulations of the point estimates of a three-part spline in g(.) as in the main text. The simulated data is the actual data for all parameters except for whether the house sold within 13 months, which is created as simulated data using an assumed value for g(.), here the baseline estimate, and then adding noise to the first stage relative markup that is independent of the instrument and normally distributed with mean zero and the indicated standard deviation. The simulation procedure is described in detail in the Appendix text.

Table 34 shows results with a normally distributed $\zeta_{h\ell t}$ that is independent of $f(z_{h\ell t})$. The assumed $g(\cdot)$ is the third-order polynomial estimate of $g(\cdot)$ shown in Figure 5. Increasing the standard deviation of $\zeta_{h\ell t}$ leads to a $g(\cdot)$ that is steeper and more linear than the baseline estimates. Other sources of variation in the relative markup that are independent of the instrument would thus likely lead to an under-estimate of the true degree of concavity.

Spurious concavity is, however, a possibility if the variance of $\zeta_{h\ell t}$ is correlated with $z_{h\ell t}$. Specifically, consider the case where the instrument captures most of the variation in the relative markup for sellers with low appreciation since purchase but little of the variation with high appreciation since purchase. Then the observed probability of sale at low $p_{h\ell t} - \tilde{p}_{h\ell t}$ would be an average of the probabilities of sale at true $p_{h\ell t} - \tilde{p}_{h\ell t}$s that are scrambled, yielding an attenuated slope for low $p_{h\ell t} - \tilde{p}_{h\ell t}$. However, at high $p_{h\ell t} - \tilde{p}_{h\ell t}$, the observed $p_{h\ell t} - \tilde{p}_{h\ell t}$ would be close to the true $p_{h\ell t} - \tilde{p}_{h\ell t}$, yielding the true slope.

Table 35 illustrates that this type of bias could cause spurious concavity, but that generating the amount of concavity I observe in the data would require an extreme amount of unobserved variation in the relative markup at low levels of appreciation since purchase. To show this, I assume the true $g(\cdot)$ is linear and let the standard deviation of $\zeta_{h\ell t}$ depend on $f(z_{h\ell t})$ as indicated in the first two rows of the table. The first column shows estimates with no noise, which are approximately

Table 35: Monte Carlo Simulations: Other Sources of Markup Variation Corr With Instrument

SD $f(z) < .01$	0	0.05	0.5	1.0	0.5
SD $f(z) \geq .01$	0	0	0	0	0.2
Lowest Tercile	-1.267	-1.256	-0.557	-.194	-0.566
	(0.382)	(0.381)	(0.414)	(.405)	(0.404)
Middle Tercile	-1.556	-1.629	-0.768	-.399	-0.786
	(1.008)	(1.000)	(1.084)	(1.070)	(1.044)
Highest Terclie	-1.321	-1.315	-1.454	-1.503	-1.234
	(0.263)	(0.257)	(0.281)	(0.275)	(0.271)
High - Low	-0.054	-0.060	-0.897	-1.310	-.668
	(0.450)	(0.436)	(.479)	(0.466)	(0.454)

Notes: Each column shows the mean and standard deviation over 1,000 Monte Carlo simulations of the point estimates of a three-part spline in g(.) as in the main text. The simulated data is the actual data for all parameters except for whether the house sold within 13 months, which is created as simulated data using an assumed value for g(.), here the baseline estimate, and then adding noise to the first stage relative markup. Here the variance of the noise depends on $f\left(z_{h\ell t}\right)$ (the estimated log relative markup) and thus the instrument. Specifically, the noise is normally distributed with a standard deviation equal to the first row if $f\left(z_{h\ell t}\right) < .01$ and the second row if $f\left(z_{h\ell t}\right) \geq .01$. This makes the noise larger for homes with more appreciation since purchase, creating the potential spurious concavity from heterskedasticity described in the text. The simulation procedure is described in detail in the Appendix text.

linear. To generate substantial spurious concavity, the instrument must be near perfect for low appreciation since sale and that the other sources of variation must have a standard deviation of over half a log point for low appreciation since purchase. In other words, the amount of other sources of variation in the relative markup must be near zero for low appreciation since purchase and over 25 times as great as the variation induced by the instrument for high appreciation since purchase.

A.4 Model

A.4.1 Laws of Motion and Value Functions

This Appendix derives the probabilities of sale and value functions for a model with a general price distribution. It then provides the laws of motion and value functions not provided in the main text.

The *ex-ante* probability of sale for a seller posting price p_t with a price distribution of Ω_t and a functional market tightness of $\tilde{\theta}_t$ is:

$$d\left(p_t, \Omega_t, \tilde{\theta}_t\right) = q\left(\tilde{\theta}_t\right)\left(1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right)\left(1 - F\left(\varepsilon_t^*\right)\right).$$

The functional market tightness is $\tilde{\theta}_t \equiv \frac{B_t}{S_t^{visited}} = \frac{\theta_t}{E_\Omega[1-G(p_t-E_\Omega[p_t]-\mu)]}$, where $\theta_t = \frac{B_t}{S_t}$ is the buyer-to-seller ratio. The probability of purchase for a buyer is the probability of a match times an integral over all sellers with the integrand equal to the probability of meeting any given seller times the probability of purchase from that seller:

$$\Pr\left[\text{Buy}|\Omega_t, \tilde{\theta}_t\right] = \frac{q\left(\tilde{\theta}_t\right)}{\tilde{\theta}_t} \int \frac{1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)}{E_\Omega\left[1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right]} \left(1 - F\left(\varepsilon_t^*\left(p_t\right)\right)\right) d\Omega\left(p_t\right)$$

Note the probability distribution the buyer integrates over is the distribution of homes that satisfy $p_t - \eta_{h,t} - E\left[p_t\right] \leq \mu$ but $\Omega\left(p_t\right)$ applies to all posted prices. Multiplying by $\frac{1-G(p_t-E_\Omega[p_t]-\mu)}{E_\Omega[1-G(p_t-E_\Omega[p_t]-\mu)]}$ converts the density $\omega\left(p_t\right)$ from the distribution of posted prices to the distribution of non-overpriced posted prices. This probability can be simplified by noting that the $E_\Omega\left[1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right]$ in the denominator of $\tilde{\theta}_t$ cancels with the same term in the integral:

$$\begin{aligned}
\Pr\left[\text{Buy}\right] &= \frac{q\left(\tilde{\theta}_t\right)}{\theta_t} \int \left(1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right) \left(1 - F\left(\varepsilon_t^*\left(p_t\right)\right)\right) d\Omega\left(p_t\right) \\
&= \frac{1}{\theta_t} \int d\left(p_t, \Omega_t, \tilde{\theta}_t\right) d\Omega\left(p_t\right) \\
&= \frac{1}{\theta_t} E_\Omega\left[d\left(p_t, \Omega_t, \tilde{\theta}_t\right)\right].
\end{aligned}$$

With Cobb-Douglas matching, things simplify because,

$$q\left(\tilde{\theta}_t\right) = \xi \tilde{\theta}_t^\gamma = \xi \frac{\theta_t^\gamma}{E_\Omega\left[\left(1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right)\right]^\gamma}.$$

The seller's value function is:

$$V_t^s = s + \beta E_t V_{t+1}^s + \max_{p_t} \left\{d\left(p_t, \Omega_t, \theta_t\right)\left(p_t - s - \beta E_t V_{t+1}^s\right)\right\},$$

where p_t is own price. For a staggered model, in period $\tau = 0$ when the seller is choosing price, the value function is:

$$V_t^{s,0} = s + \beta E_t V_{t+1}^{s,1} + \max_{p_t^0} \left\{d\left(p_t^0, \Omega_t, \theta_t\right)\left(p_t^0 - s - \beta E_t V_{t+1}^{s,1}\right)\right\},$$

and in period $\tau \neq 0$ when price is fixed, the value function is:

$$V_t^{s,\tau} = s + \beta E_t V_{t+1}^{s,\tau+1} + d\left(p_t^\tau, \Omega_t, \theta_t\right)\left(p_t^\tau - s - \beta E_t V_{t+1}^{s,\tau+1}\right).$$

$p_t^\tau = p_{t-1}^{\tau-1}$ and $V_t^{s,N} = V_t^{s,0}$, so that $V_{t+1}^{s,\tau+1}$ corresponds to $V_{t+1}^{s,0}$ when $\tau = N-1$.

The buyer's value function integrates out over each type of seller, multiplying the probability of purchase from that type of seller by the surplus obtained from purchasing from that type of seller:

$$V_t^b = b + \beta E_t V_{t+1}^b + \frac{1}{\tilde{\theta}_t}\int d\left(p_t, \Omega_t, \tilde{\theta}_t\right)\left(\begin{array}{l} V_t^h - p_t^\tau - b - \beta V_{t+1}^b + \varepsilon_t^{*,\tau}\left(p_t^\tau\right) \\ +E\left[\varepsilon - \varepsilon_t^{*,\tau}\left(p_t^\tau\right) | \varepsilon > \varepsilon_t^{*,\tau}\left(p_t^\tau\right)\right] \end{array}\right) d\Omega\left(p_t\right).$$

The purchase decisions from each type of seller are optimal, so

$$\varepsilon_t^{*,\tau}\left(p_t^\tau\right) = p_t^\tau + b + \beta V_{t+1}^b - V_t^h.$$

Plugging this into V_t^b gives:

$$\begin{aligned} V_t^b &= b + \beta E_t V_{t+1}^b + \frac{1}{\theta_t}\int d\left(p_t, \Omega_t, \tilde{\theta}_t\right) E\left[\varepsilon - \varepsilon_t^{*,\tau}\left(p_t^\tau\right) | \varepsilon > \varepsilon_t^{*,\tau}\left(p_t^\tau\right)\right] d\Omega\left(p_t\right) \\ &= b + \beta E_t V_{t+1}^b + \frac{1}{\chi\theta_t}\int d\left(p_t, \Omega_t, \tilde{\theta}_t\right) d\Omega\left(p_t\right). \end{aligned}$$

Given the setup presented in the main text and summarized in Figure 6, the laws of motion for an arbitrary list price distribution Ω_t are:

$$B_t = \left(1 - \frac{1}{\theta_{t-1}} E_{\Omega_{t-1}}\left[d\left(p_{t-1}, \Omega_{t-1}, \tilde{\theta}_{t-1}\right)\right]\right) B_{t-1} \tag{48}$$

$$+\lambda^r\left(1 - K\left(k_{t-1}^*\right)\right) R_{t-1} + \left(1 - LK\left(k_{t-1}^*\right)\right)\lambda^h\left(1 - C\left(c_{t-1}^*\right)\right) H_{t-1} \tag{49}$$

$$S = \left(1 - E_{\Omega_{t-1}}\left[d\left(p_{t-1}, \Omega_{t-1}, \tilde{\theta}_{t-1}\right)\right]\right) S_{t-1} + \lambda^h\left(1 - C\left(c_{t-1}^*\right)\right) H_{t-1} \tag{50}$$

$$H_t = E_{\Omega_{t-1}}\left[d\left(p_{t-1}^\tau, \Omega_{t-1}, \tilde{\theta}_{t-1}\right)\right] S_{t-1} + \left(1 - \lambda^h\left(1 - C\left(c_{t-1}^*\right)\right)\right) H_{t-1} \tag{51}$$

$$R_t = \left(1 - \lambda^r\left(1 - K\left(k_{t-1}^*\right)\right)\right) R_{t-1} + LK\left(k_{t-1}^*\right)\lambda^h\left(1 - C\left(c_{t-1}^*\right)\right) H_{t-1}. \tag{52}$$

For staggered pricing there are laws of motion for each S^{τ}:

$$S_t^{\tau} = \left(1 - d\left(p_{t-1}^{\tau-1}, \Omega_{t-1}, \tilde{\theta}_{t-1}\right)\right) S_{t-1}^{\tau-1} \; \forall \; \tau \neq 0 \tag{53}$$

$$S_t^{0} = \left(1 - d\left(p_{t-1}^{\tau-1}, \Omega_{t-1}, \tilde{\theta}_{t-1}\right)\right) S_{t-1}^{N-1} + \lambda^h \left(1 - C\left(c_{t-1}^*\right)\right) H_{t-1}. \tag{54}$$

The value functions for homeowners and renters are,

$$V_t^h = h + \beta \begin{bmatrix} \lambda^h \left(1 - C\left(c_{t+1}^*\right)\right) E_t \left[V_{t+1}^s + LV^0 + (1-L) V_{t+1}^b\right] \\ -\lambda^h C\left(c_{t+1}^*\right) E\left[c | c < c_t^*\right] + \left(1 - \lambda^h \left(1 - C\left(c_{t+1}^*\right)\right)\right) E_t V_{t+1}^h \end{bmatrix} \tag{55}$$

$$V_t^r = u + x_t + \beta \begin{bmatrix} \lambda^r \left(1 - K\left(k_{t+1}^*\right)\right) E_t V_{t+1}^b - \lambda^r K\left(k_{t+1}^*\right) E\left[k | k < k_t^*\right] \\ + \left(1 - \lambda^r \left(1 - K\left(k_{t+1}^*\right)\right)\right) E_t V_{t+1}^r \end{bmatrix}, \tag{56}$$

where x_t is the stochastic AR(1) shock to the flow utility of renting. In the staggered model V_{t+1}^s is replaced by $V_{t+1}^{s,0}$. The value functions are standard with the exception of terms for the expected cost paid by a homeowner and renter if they decide not to enter the housing market, $E\left[c | c < c_t^*\right]$ and $E\left[k | k < k_t^*\right]$, respectively.

A.4.2 Proofs

Lemma 3: Optimal Flexible Price Setting The seller's value function is:

$$V_t^s = s + \beta V_{t+1}^s + \max_p \left\{ d\left(p_t, \Omega_t, \tilde{\theta}_t\right) \left(p_t - s - \beta V_{t+1}^s\right) \right\}.$$

The seller's optimal price is defined by the first order condition:

$$p_t = s + \beta V_{t+1}^s - \frac{d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\frac{\partial d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\partial p_t}},$$

where $d\left(p_t, \Omega_t, \tilde{\theta}_t\right)$ is defined in equation (12) and $\varepsilon_t^* = b + \beta V_{t+1}^b + p_t - V_t^h$. The impact of an individual's price on $E_\Omega\left[1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right]$ is infinitesimal so,

$$\frac{\partial d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\partial p_t} = -d\left(p_t, \Omega_t, \tilde{\theta}_t\right) \left[\frac{f\left(\varepsilon_t^*\right)}{1 - F\left(\varepsilon_t^*\right)} + \frac{g\left(\left(p_t - E_\Omega\left[p_t\right] - \mu\right)\right)}{1 - G\left(p_t - E_\Omega\left[p_t\right] - \mu\right)}\right],$$

which gives:

$$\begin{aligned} p &= s + \beta V_{t+1}^{s} + \frac{1}{\frac{f(\varepsilon_t^*)}{1-F(\varepsilon_t^*)} + \frac{g((p_t - E_\Omega[p_t] - \mu))}{1-G(p_t - E_\Omega[p_t] - \mu)}} \\ &= s + \beta V_{t+1}^{s} + \frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p_t - E_\Omega[p_t] - \mu}{\sigma}\right)} + \chi}. \end{aligned}$$

For uniqueness, the second order condition is:

$$\frac{\partial^2 d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\partial p_t^2}\left(p_t - s - \beta V_{t+1}^{s}\right) + 2\frac{\partial d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\partial p_t} < 0,$$

or equivalently,

$$\frac{\partial^2 d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\partial p_t^2} < \frac{2}{d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}\left[\frac{\partial d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\partial p_t}\right]^2.$$

This condition holds locally, as $\frac{\partial^2 d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\partial p_t^2} < 0$ around the equilibrium price. However, $d\left(p_t, \Omega_t, \tilde{\theta}_t\right)$ is not globally concave as illustrated by Figure 7. Because it is convex as $p_t - E_\Omega[p_t] \to \infty$, the Lemma specifies that the optimal price is unique on an interval bounded away from $p_t = \infty$. Intuitively, there may be an equilibrium with close to no trade due to the non-concavity, but this equilibrium is assumed away. Numerical simulations show that the local optimum is the global optimum for all parameter values considered.

Lemma 4: Optimal Staggered Price Setting The price-setting seller's value function is:

$$V_t^{s,0} = \max_p \left\{ s + \beta V_{t+1}^{s,1}(p) + d\left(p, \Omega_t, \tilde{\theta}_t\right)\left(p - s - \beta V_{t+1}^{s,1}(p)\right)\right\},$$

where

$$V_t^{s,\tau}(p) = s + \beta V_{t+1}^{s,\tau+1}(p) + d\left(p, \Omega_t, \tilde{\theta}_t\right)\left(p - s - \beta V_{t+1}^{s,\tau+1}(p)\right),$$

and $V_t^N = V_t^0$. The first order condition is:

$$\beta\left(1 - d\left(p, \Omega_t, \tilde{\theta}_t\right)\right) E_t \frac{\partial V_{t+1}^{s,1}}{\partial p} + d\left(p, \Omega_t, \tilde{\theta}_t\right) E_t\left[1 + \frac{\frac{\partial d\left(p, \Omega_t, \tilde{\theta}_t\right)}{\partial p}}{d\left(p, \Omega_t, \tilde{\theta}_t\right)}\left(p - s - \beta V_{t+1}^{s,1}(p)\right)\right] = 0,$$

where for $\tau < N-1$,

$$\begin{aligned} E_t \frac{\partial V_t^{s,\tau}(p)}{\partial p} &= \beta\left(1 - d\left(p, \Omega_t, \tilde{\theta}_t\right)\right) E_t \frac{\partial V_{t+1}^{s,\tau+1}(p)}{\partial p} \\ &\quad + d\left(p, \Omega_t, \tilde{\theta}_t\right) E_t \left[1 + \frac{\frac{\partial d\left(p,\Omega_t,\tilde{\theta}_t\right)}{\partial p}}{d\left(p, \Omega_t, \tilde{\theta}_t\right)} \left(p - s - \beta V_{t+1}^{s,\tau+1}(p)\right)\right], \end{aligned}$$

and,

$$E_t \frac{\partial V_t^{s,N-1}(p)}{\partial p} = d\left(p, \Omega_t, \tilde{\theta}_t\right) E_t \left[1 + \frac{\frac{\partial d\left(p,\Omega_t,\tilde{\theta}_t\right)}{\partial p}}{d\left(p, \Omega_t, \tilde{\theta}_t\right)} \left(p - s - \beta_{t+1}^{s,0}\right)\right].$$

Defining $D_t^j(p) = E_t\left[\prod_{\tau=0}^{j-1}\left(1 - d^\tau\left(p, \Omega_{t+\tau}, \tilde{\theta}_{t+\tau}\right)\right)\right] d\left(p, \Omega_{t+j}, \tilde{\theta}_{t+j}\right)$ and substituting $\frac{\partial V_{t+1}^{s,1}}{\partial p}, ..., \frac{\partial V_{t+N-1}^{s,N-1}}{\partial p}$ into the first order condition gives:

$$\sum_{\tau=0}^{N-1} \beta^\tau D_t^\tau(p) E_t \left[1 + \frac{\frac{\partial d\left(p,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)}{\partial p}}{d\left(p, \Omega_{t+\tau}, \tilde{\theta}_{t+\tau}\right)} \left(p - s - \beta V_{t+\tau+1}^{s,\tau+1}\right)\right] = 0.$$

Rearranging gives:

$$p = \frac{\sum_{\tau=0}^{N-1} \beta^\tau D_t^\tau(p) E_t \left[1 + \frac{-\frac{\partial d\left(p,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)}{\partial p}}{d\left(p,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)} \left(s + \beta V_{t+\tau+1}^{s,\tau+1}\right)\right]}{\sum_{\tau=0}^{N-1} \beta^\tau D_t^\tau(p) E_t \left[\frac{-\frac{\partial d\left(p,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)}{\partial p}}{d\left(p,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)}\right]},$$

which, defining $\Psi_t^\tau = E_t\left[\frac{-\frac{\partial d\left(p_t,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)}{\partial p_t}}{d\left(p_t,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)}\right]$ and $\varphi_t^\tau = s + E_t V_{t+\tau+1}^{s,\tau+1} + \frac{1}{\Psi_t^\tau}$, simplifies to,

$$p = \frac{\sum_{\tau=0}^{N-1} \beta^\tau D_t^\tau(p) \Psi_t^\tau \varphi_t^\tau}{\sum_{\tau=0}^{N-1} \beta^\tau D_t^\tau(p) \Psi_t^\tau}.$$

For uniqueness, that the second order condition is:

$$E_t \left\{ \begin{array}{l} \beta\left(1 - d\left(p, \Omega_t, \tilde{\theta}_t\right)\right) \frac{\partial^2 V_{t+1}^{s,1}}{\partial p^2} - 2\beta \frac{\partial d\left(p,\Omega_t,\tilde{\theta}_t\right)}{\partial p} \frac{\partial V_{t+1}^{s,1}}{\partial p} \\ +2\frac{\partial d\left(p,\Omega_t,\tilde{\theta}_t\right)}{\partial p} + \frac{\partial^2 d\left(p,\Omega_t,\tilde{\theta}_t\right)}{\partial p^2}\left(p - s - \beta V_{t+1}^{s,1}(p)\right) \end{array} \right\} < 0,$$

where

$$
\begin{aligned}
E_t \frac{\partial V_t^{s,\tau}(p)}{\partial p} &= E_t \left\{ \begin{array}{c} \beta \left(1 - d\left(p, \Omega_t, \tilde{\theta}_t\right)\right) \frac{\partial V_{t+1}^{s,\tau+1}(p)}{\partial p} \\ +d\left(p, \Omega_t, \tilde{\theta}_t\right) + \frac{\partial d\left(p, \Omega_t, \tilde{\theta}_t\right)}{\partial p} \left(p - s - \beta V_{t+1}^{s,\tau+1}(p)\right) \end{array} \right\}, \\
E_t \frac{\partial^2 V_t^{s,\tau}(p)}{\partial p^2} &= E_t \left\{ \begin{array}{c} \beta \left(1 - d\left(p, \Omega_t, \tilde{\theta}_t\right)\right) \frac{\partial^2 V_{t+1}^{s,\tau+1}(p)}{\partial p^2} - 2\beta \frac{\partial d\left(p, \Omega_t, \tilde{\theta}_t\right)}{\partial p} \frac{\partial V_{t+1}^{s,\tau+1}(p)}{\partial p} \\ +2\frac{\partial d\left(p, \Omega_t, \tilde{\theta}_t\right)}{\partial p} + \frac{\partial^2 d\left(p, \Omega_t, \tilde{\theta}_t\right)}{\partial p^2} \left(p - s - \beta V_{t+1}^{s,\tau+1}(p)\right) \end{array} \right\},
\end{aligned}
$$

and

$$
\begin{aligned}
E_t \frac{\partial V_t^{s,N-1}(p)}{\partial p} &= E_t \left[d\left(p, \Omega_t, \tilde{\theta}_t\right) + \frac{\partial d\left(p, \Omega_t, \tilde{\theta}_t\right)}{\partial p} \left(p - s - \beta_{t+1}^{s,0}\right) \right] \\
E_t \frac{\partial^2 V_t^{s,N-1}(p)}{\partial p^2} &= E_t \left[2\frac{\partial d\left(p, \Omega_t, \tilde{\theta}_t\right)}{\partial p} + \frac{\partial^2 d\left(p, \Omega_t, \tilde{\theta}_t\right)}{\partial p^2} \left(p - s - \beta_{t+1}^{s,0}\right) \right].
\end{aligned}
$$

The second order condition can be rewritten as:

$$
\sum_{\tau=0}^{N-1} \beta^\tau \prod_{j=0}^{\tau} \left(1 - d\left(p, \Omega_{t+j}, \tilde{\theta}_{t+j}\right)\right) \left\{ \begin{array}{c} 2\frac{\partial d\left(p, \Omega_{t+\tau}, \tilde{\theta}_{t+\tau}\right)}{\partial p} + \frac{\partial^2 d\left(p, \Omega_{t+\tau}, \tilde{\theta}_{t+\tau}\right)}{\partial p^2} \left(p - s - \beta V_{t+\tau+1}^{s,\tau+1}(p)\right) \\ -2\beta \frac{\partial d\left(p, \Omega_{t+\tau}, \tilde{\theta}_{t+\tau}\right)}{\partial p} \frac{\partial V_{t+\tau+1}^{s,\tau+1}}{\partial p} \end{array} \right\} < 0,
$$

where,

$$
E_t \frac{\partial V_{t+\tau}^{s,t+\tau}}{\partial p} = E_t \sum_{k=t+\tau}^{t+N-1} \left\{ \begin{array}{c} \beta^{k-t-\tau} \left[\prod_{j=t+\tau}^{k} \left(1 - d\left(p, \Omega_j, \tilde{\theta}_j\right)\right) \right] d\left(p, \Omega_{t+\tau+k}, \tilde{\theta}_{t+\tau+k}\right) \\ \times \left[d\left(p, \Omega_{t+\tau+k}, \tilde{\theta}_{t+\tau+k}\right) + \frac{\partial d\left(p, \Omega_{t+\tau+k}, \tilde{\theta}_{t+\tau+k}\right)}{\partial p} \left(p - s - \beta V_{t+\tau+k+1}^{s,t+\tau+k+1}(p)\right) \right] \end{array} \right\}.
$$

Combining gives

$$
\begin{aligned}
& E_t \sum_{\tau=0}^{N-1} \beta^\tau \prod_{j=0}^{\tau} \left(1 - d\left(p, \Omega_{t+j}, \tilde{\theta}_{t+j}\right)\right) \left\{ \begin{array}{c} 2\frac{\partial d\left(p,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)}{\partial p} + \\ \frac{\partial^2 d\left(p,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)}{\partial p^2}\left(p - s - \beta V_{t+\tau+1}^{s,\tau+1}(p)\right) \end{array} \right\} \\
& -E_t \sum_{\tau=0}^{N-1} \beta^\tau \left\{ \begin{array}{c} 2\beta \frac{\partial d\left(p,\Omega_{t+\tau},\tilde{\theta}_{t+\tau}\right)}{\partial p} \sum_{k=t+\tau}^{t+N-1} \beta^{k-t-\tau} \left[\prod_{j=\tau}^{k}\left(1 - d\left(p, \Omega_j, \tilde{\theta}_j\right)\right)\right] \\ \times \frac{\partial d\left(p,\Omega_{t+\tau+k},\tilde{\theta}_{t+\tau+k}\right)}{\partial p}\left(p - s - \beta V_{t+\tau+k+1}^{s,t+\tau+k+1}(p)\right) \end{array} \right\} \\
< \quad & E_t \sum_{\tau=0}^{N-1} \beta^\tau \left\{ 2\beta \frac{\partial d\left(p, \Omega_{t+\tau}, \tilde{\theta}_{t+\tau}\right)}{\partial p} \sum_{k=t+\tau}^{t+N-1} \beta^{k-t-\tau} \left[\prod_{j=\tau}^{k}\left(1 - d\left(p, \Omega_j, \tilde{\theta}_j\right)\right)\right] d\left(p, \Omega_{t+\tau+k}, \tilde{\theta}_{t+\tau+k}\right) \right\}.
\end{aligned}
$$

This condition must hold for uniquness. As with the flexible case, $d\left(p_t, \Omega_t, \tilde{\theta}_t\right)$ is concave locally but not globally, as illustrated by Figure 7. Because it is convex as $p_t - E_\Omega[p_t] \to \infty$, the Lemma specifies the optimal price is unique on an interval bounded away from $p_t = \infty$. Locally, the left hand side of the final condition is negative, while the right hand side is indeterminate. In all simulations considered, this condition holds, and numerical simulations show that the local optimum is the global optimum for all parameter values considered.

A.4.3 Frictionless Model

In the frictionless case ($N = 1$ for the staggered model, $\alpha = 1$ for the backward-looking model), the definition of an equilibrium simplifies to:

Definition 7 *Equilibrium with flexible pricing is defined by a price p_t, purchase cutoffs ε_t^*, and seller, buyer, homeowner, and renter value functions V_t^s, V_t^b, V_t^h, and V_t^r, entry cutoffs c_t^* and k_t^*, and stocks of each type of agent B_t, S_t, H_t, and R_t satisfying:*

1. *Optimal pricing (15). All sellers set the same price so $p_t = E_\Omega[p_t]$*
2. *Optimal purchasing decisions by buyers: $\varepsilon_t^* = p_t + b + \beta V_{t+1}^b - V_t^h$*
3. *Demand curve (12) simplifies to $q(\theta_t)(1 - F(\varepsilon_t^*))$*
4. *Optimal entry decisions by homeowners and renters who receive shocks (9) and (10)*

5. *The value functions for buyers (13), sellers (14), renters (56), and homeowners (55)*

6. *The laws of motion for all agents (48), (51), (52), and (50).*

For the reader's convenience, the simplified general equilibrium system without a shock is reproduced below. Note that $d\left(p_t, \Omega_t, \tilde{\theta}_t\right) = \xi\theta^{\gamma}\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon_t^*\right)\right)$ is plugged in everywhere:

$$
\begin{aligned}
B_t &= \left(1 - \frac{1}{\theta_{t-1}}\xi\theta_{t-1}^{\gamma}\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon_t^*\right)\right)\right) B_{t-1} \\
&\quad + \lambda^r\left(1 - K\left(k_{t-1}^*\right)\right) R_{t-1} + \left(1 - LK\left(k^*\right)\right)\lambda^h\left(1 - C\left(c_{t-1}^*\right)\right) H_{t-1} \\
S_t &= \left(1 - \xi\theta_{t-1}^{\gamma}\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon_t^*\right)\right)\right) S_{t-1} + \lambda^h\left(1 - C\left(c_{t-1}^*\right)\right) H_{t-1} \\
H_t &= \xi\theta_{t-1}^{\gamma}\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon_t^*\right)\right) S_{t-1} + \left(1 - \lambda^h\left(1 - C\left(c_{t-1}^*\right)\right)\right) H_{t-1} \\
R_t &= \left(1 - \lambda^r\left(1 - K\left(k_{t-1}^*\right)\right)\right) R_{t-1} + LK\left(k^*\right)\lambda^h\left(1 - C\left(c_{t-1}^*\right)\right) H_{t-1} \\
V_t^h &= h + \beta\begin{bmatrix} \lambda^h\left(1 - C\left(c_{t+1}^*\right)\right)\left[V_{t+1}^s + LV^0 + \left(1 - L\right)V_{t+1}^b\right] \\ -\lambda^h C\left(c_{t+1}^*\right) E\left[c | c < c_{t+1}^*\right] + \left(1 - \lambda^h\left(1 - C\left(c_{t+1}^*\right)\right)\right) V_{t+1}^h \end{bmatrix} \\
V_t^r &= u + x_t + \beta\left[\lambda^r\left(1 - K\left(k_{t+1}^*\right)\right) V_{t+1}^b - \lambda^r K\left(k_{t+1}^*\right) E\left[k | k < k_{t+1}^*\right] + \left(1 - \lambda^r\left(1 - K\left(k_{t+1}^*\right)\right)\right) V_{t+1}^r\right] \\
V_t^b &= b + \beta E_t V_{t+1}^b + \frac{\xi\theta_t^{\gamma}\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon_t^*\right)\right)}{\chi\theta_t} \\
V_t^s &= s + \beta E_t V_{t+1}^s + \xi\theta_t^{\gamma}\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon_t^*\right)\right)\frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(\frac{\mu}{\sigma}\right)} + \chi} \\
\varepsilon_t^* &= b + \beta E_t V_{t+1}^b + p_t^i - V_t^h \\
p_t &= s + \beta E_t V_{t+1}^s + \frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(\frac{\mu}{\sigma}\right)} + \chi} \\
c_t^* &= V_t^h - \left(V_t^s + \left(1 - L\right)V_t^b + LV_t^0\right) \\
k_t^* &= V_t^r - V_t^b \\
x_t &= \rho x_{t+1} + \eta,\ \eta \sim N\left(0, \sigma_{\eta}^2\right)
\end{aligned}
$$

A.4.4 Backward-Looking Model

The equilibrium of the staggered pricing model is defined by Definition 6. For the reader's convenience, the general equilibrium system is reproduced below.

Full Backward-Looking Model

$$
\begin{aligned}
B_t &= \left(1-\frac{1}{\theta_{t-1}}\left[\alpha d\left(p_{t-1}^N,\Omega_{t-1},\tilde{\theta}_{t-1}\right)+(1-\alpha)\,d\left(p_{t-1}^R,\Omega_{t-1},\tilde{\theta}_{t-1}\right)\right]\right)B_{t-1}\\
&\quad +\lambda^r\left(1-K\left(k_{t-1}^*\right)\right)R_{t-1}+\left(1-LK\left(k_{t-1}^*\right)\right)\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)H_{t-1}\\
S_t &= \left(1-\alpha d\left(p_{t-1}^N,\Omega_{t-1},\tilde{\theta}_{t-1}\right)-(1-\alpha)\,d\left(p_{t-1}^R,\Omega_{t-1},\tilde{\theta}_{t-1}\right)\right)S_{t-1}^1+\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)H_{t-1}\\
H_t &= d\left(p_{t-1}^N,\Omega_{t-1},\tilde{\theta}_{t-1}\right)\alpha S_{t-1}+d\left(p_{t-1}^R,\Omega_{t-1},\tilde{\theta}_{t-1}\right)(1-\alpha)\,S_{t-1}+\left(1-\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)\right)H_{t-1}\\
R_t &= \left(1-\lambda^r\left(1-K\left(k^*\right)\right)\right)R_{t-1}+LK\left(k_{t-1}^*\right)\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)H_{t-1}\\
V_t^h &= h+\beta\begin{bmatrix}\lambda^h\left(1-C\left(c_{t+1}^*\right)\right)\left[V_{t+1}^{s,R}+LV^0+(1-L)\,V_{t+1}^b\right]\\ -\lambda^h C\left(c_{t+1}^*\right)E\left[c|c<c_{t+1}^*\right]+\left(1-\lambda^h\left(1-C\left(c_{t+1}^*\right)\right)\right)V_{t+1}^h\end{bmatrix}\\
V_t^r &= u+x_t+\beta\left[\lambda^r\left(1-K\left(k_{t+1}^*\right)\right)V_{t+1}^b-\lambda^r K\left(k_{t+1}^*\right)E\left[k|k<k_{t+1}^*\right]+\left(1-\lambda^r\left(1-K\left(k_{t+1}^*\right)\right)\right)V_{t+1}^r\right]\\
V_t^b &= b+\beta E_t V_{t+1}^b+\frac{1}{\chi\theta_t}\left[\alpha d\left(p_t^N,\Omega_t,\tilde{\theta}_t\right)+(1-\alpha)\,d\left(p_t^R,\Omega_t,\tilde{\theta}_t\right)\right]\\
V_t^{s,R} &= s+\beta E_t V_{t+1}^{s,R}+d\left(p_t^R,\Omega_t,\tilde{\theta}_t\right)\frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p-E[p]-\mu}{\sigma}\right)}+\chi}\\
\varepsilon_t^{i*} &= b+\beta E_t V_{t+1}^b+p_t^i-V_t^h\ \forall\ i\\
d\left(p_t^i,\Omega_t,\theta\right) &= q\left(\theta_t\right)\left(1-F\left(\varepsilon_t^{i*}\right)\right)\frac{\frac{\exp\left(-\frac{p_t^i-E[p]-\mu}{\sigma}\right)}{1+\exp\left(-\frac{p_t^i-E[p]-\mu}{\sigma}\right)}}{\left[\alpha\frac{\exp\left(-\frac{p_t^N-E[p]-\mu}{\sigma}\right)}{1+\exp\left(-\frac{p_t^N-E[p]-\mu}{\sigma}\right)}+(1-\alpha)\frac{\exp\left(-\frac{p_t^R-E[p]-\mu}{\sigma}\right)}{1+\exp\left(-\frac{p_t^R-E[p]-\mu}{\sigma}\right)}\right]^\gamma}\\
c_t^* &= V_t^h-\left(V_t^s+(1-L)\,V_t^b+LV_t^0\right)\\
k_t^* &= V_t^r-V_t^b\\
p_t^N &= \frac{p_{t-2}+p_{t-3}+p_{t-4}}{3}+\phi\left(\frac{p_{t-2}+p_{t-3}+p_{t-4}}{3}-\frac{p_{t-5}+p_{t-6}+p_{t-7}}{3}\right)\\
p_t^R &= s+\beta E_t V_{t+1}^s+\frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p-E[p]-\mu}{\sigma}\right)}+\chi}\\
E\left[p_t\right] &= \alpha p_t^N+(1-\alpha)\,p_t^R\\
p_t &= \frac{\alpha d\left(p_t^N,\Omega_t,\tilde{\theta}_t\right)p^N+(1-\alpha)\,d\left(p_t^R,\Omega_t,\tilde{\theta}_t\right)p_t^R}{\alpha d\left(p_t^N,\Omega_t,\tilde{\theta}_t\right)+(1-\alpha)\,d\left(p_t^R,\Omega_t,\tilde{\theta}_t\right)}\\
x_t &= \rho x_{t+1}+\eta,\ \eta\sim N\left(0,\sigma_\eta^2\right)
\end{aligned}
$$

A.4.5 Staggered Pricing Model

The equilibrium of the staggered pricing model is defined by Definition 5. For the reader's convenience, the general equilibrium system for two alternating groups of sellers used for most of the simulations without a shock is reproduced below.

$$
\begin{aligned}
B_t &= \left(1-\frac{1}{\theta_{t-1}}\left[\frac{S_{t-1}^0}{S_{t-1}}d\left(p_{t-1}^0,\Omega_{t-1},\tilde{\theta}_{t-1}\right)+\frac{S_{t-1}^1}{S_{t-1}}d\left(p_{t-1}^1,\Omega_{t-1},\tilde{\theta}_{t-1}\right)\right]\right)B_{t-1}\\
&\quad +\lambda^r\left(1-K\left(k_{t-1}^*\right)\right)R_{t-1}+\left(1-LK\left(k_{t-1}^*\right)\right)\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)H_{t-1}\\
S_t^0 &= \left(1-d\left(p_{t-1}^1,\Omega_{t-1},\tilde{\theta}_{t-1}\right)\right)S_{t-1}^1+\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)H_{t-1}\\
S_t^1 &= \left(1-d\left(p_{t-1}^0,\Omega_{t-1},\tilde{\theta}_{t-1}\right)\right)S_{t-1}^0\\
S_t &= S_t^0+S_t^1\\
H_t &= d\left(p_{t-1}^0,\Omega_{t-1},\tilde{\theta}_{t-1}\right)S_{t-1}^0+d\left(p_{t-1}^1,\Omega_{t-1},\tilde{\theta}_{t-1}\right)S_{t-1}^1+\left(1-\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)\right)H_{t-1}\\
R_t &= \left(1-\lambda^r\left(1-K\left(k^*\right)\right)\right)R_{t-1}+LK\left(k_{t-1}^*\right)\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)H_{t-1}\\
V_t^h &= h+\beta\begin{bmatrix}\lambda^h\left(1-C\left(c_{t+1}^*\right)\right)\left[V_{t+1}^{s,0}+LV^0+(1-L)V_{t+1}^b\right]\\ -\lambda^h C\left(c_{t+1}^*\right)E\left[c|c<c_{t+1}^*\right]+\left(1-\lambda^h\left(1-C\left(c_{t+1}^*\right)\right)\right)V_{t+1}^h\end{bmatrix}\\
V_t^r &= u+x_t+\beta\left[\lambda^r\left(1-K\left(k_{t+1}^*\right)\right)V_{t+1}^b-\lambda^r K\left(k_{t+1}^*\right)E\left[k|k<k_{t+1}^*\right]+\left(1-\lambda^r\left(1-K\left(k_{t+1}^*\right)\right)\right)V_{t+1}^r\right]\\
V_t^b &= b+\beta E_t V_{t+1}^b+\frac{1}{\chi\theta_t}\left[\frac{S_t^0}{S_t}d\left(p_t^0,\Omega_t,\tilde{\theta}_t\right)+\frac{S_{t-1}^1}{S_{t-1}}d\left(p_t^1,\Omega_t,\tilde{\theta}_t\right)\right]\\
V_t^{s,0} &= s+\beta E_t V_{t+1}^{S,1}+d\left(p_t^0,\Omega_t,\tilde{\theta}_t\right)\left(p_t^0-s-\beta E_t V_{t+1}^{s,1}\right)\\
V_t^{s,1} &= s+\beta E_t V_{t+1}^{S,0}+d\left(p_t^1,\Omega_t,\tilde{\theta}_t\right)\left(p_t^1-s-\beta E_t V_{t+1}^{s,0}\right)\\
\varepsilon_t^{i*} &= b+\beta E_t V_{t+1}^b+p_t^i-V_t^h\ \forall\ i\\
d\left(p_t^i,\Omega_t,\theta\right) &= q\left(\theta_t\right)\left(1-F\left(\varepsilon_t^{i*}\right)\right)\frac{\frac{\exp\left(-\frac{p_t^i-E[p]-\mu}{\sigma}\right)}{1+\exp\left(-\frac{p_t^i-E[p]-\mu}{\sigma}\right)}}{\left[\frac{S_t^0}{S_t}\frac{\exp\left(-\frac{p_t^0-E[p]-\mu}{\sigma}\right)}{1+\exp\left(-\frac{p_t^0-E[p]-\mu}{\sigma}\right)}+\frac{S_t^1}{S_t}\frac{\exp\left(-\frac{p_t^1-E[p]-\mu}{\sigma}\right)}{1+\exp\left(-\frac{p_t^1-E[p]-\mu}{\sigma}\right)}\right]^\gamma}\\
E\left[p\right] &= \frac{S_t^0}{S_t}p_t^0+\frac{S_t^1}{S_t}p_t^1\\
p_t^1 &= p_{t-1}^0\\
p_t^0 &= \frac{\begin{array}{c}d\left(p_t^0,\Omega_t,\theta_t\right)\left[\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p_t^0-E[p_t]-\mu}{\sigma}\right)}+\chi\right]\left(s+\beta E_t V_{t+1}^{s,1}+\frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p_t^0-E[p_t]-\mu}{\sigma}\right)}+\chi}\right)\\ +\beta\left(1-d\left(p_t^0,\Omega_t,\theta_t\right)\right)E_t d\left(p_t^0,\Omega_{t+1},\theta_{t+1}\right)\times\\ \left[\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p_t^0-E[p_{t+1}]-\mu}{\sigma}\right)}+\chi\right]\left(s+\beta V_{t+2}^{s,0}+\frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p_t^0-E[p_{t+1}]-\mu}{\sigma}\right)}+\chi}\right)\end{array}}{\begin{array}{c}d\left(p_t^0,\Omega_t,\theta_t\right)\left[\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p_t^0-E[p_t]-\mu}{\sigma}\right)}+\chi\right]\\ +\beta\left(1-d\left(p_t^0,\Omega_t,\theta_t\right)\right)E_t d\left(p_t^0,\Omega_{t+1},\theta_{t+1}\right)\left[\frac{1}{\sigma}\frac{1}{1+\exp\left(-\frac{p_t^0-E[p_{t+1}]-\mu}{\sigma}\right)}+\chi\right]\end{array}}
\end{aligned}
$$

$$
\begin{aligned}
c_t^* &= V_t^h - \left(V_t^s + (1-L)\, V_t^b + L V_t^0\right) \\
k_t^* &= V_t^r - V_t^b \\
x_t &= \rho x_{t+1} + \eta,\ \eta \sim N\left(0, \sigma_\eta^2\right)
\end{aligned}
$$

A.4.6 Steady State

All three variants of the model share a unique steady state that can be found by equating the value of the endogenous variables across time periods. Steady state values are denoted without t subscripts. Using the fact that a fixed housing stock of mass one and a fixed population in steady state of mass Pop imply:

$$
\begin{aligned}
H + S &= 1 \\
B + H + R &= N.
\end{aligned}
$$

The laws of motion and $H + S = 1$ give:

$$
\begin{aligned}
H &= \frac{q(\theta)\left[1 - G(-\mu)\right]^{1-\gamma}\left(1 - F(\varepsilon^*)\right)}{q(\theta)\left[1 - G(-\mu)\right]^{1-\gamma}\left(1 - F(\varepsilon^*)\right) + \lambda^h\left(1 - C(c^*)\right)} \\
S &= \frac{\lambda^h\left(1 - C(c^*)\right)}{q(\theta)\left[1 - G(-\mu)\right]^{1-\gamma}\left(1 - F(\varepsilon^*)\right) + \lambda^h\left(1 - C(c^*)\right)}.
\end{aligned}
$$

The laws of motion along with $B + H + R = Pop$ imply:

$$
\begin{aligned}
R &= \frac{LK(k^*)\,\lambda^h\left(1 - C(c^*)\right)}{\lambda^r\left(1 - K(k^*)\right)} H \\
B &= \frac{\lambda^h\left(1 - C(c^*)\right)}{\frac{q(\theta)}{\theta}\left[1 - G(-\mu)\right]^{1-\gamma}\left(1 - F(\varepsilon^*)\right)} H \\
\frac{LK(k^*)}{\lambda^r\left(1 - K(k^*)\right)} &= \frac{(Pop - 1)}{\lambda^h\left(1 - C(c^*)\right)} + \frac{(Pop - \theta)}{q(\theta)\left[1 - G(-\mu)\right]^{1-\gamma}\left(1 - F(\varepsilon^*)\right)}.
\end{aligned}
$$

The steady state value functions are:

$$
\begin{aligned}
V^h &= \frac{h + \beta\lambda^h\left(1 - C(c^*)\right)\left[V^s + LV^0 + (1-L)\, V^b\right] - \beta\lambda^h C(c^*)\, E\left[c | c < c^*\right]}{1 - \beta\left(1 - \lambda^h\left(1 - C(c^*)\right)\right)} \\
V^r &= \frac{u + \beta\lambda^r\left(1 - K(k^*)\right) V^b - \beta\lambda^r K(k^*)\, E\left[k | k < k^*\right]}{1 - \beta\left(1 - \lambda^r\left(1 - K(k^*)\right)\right)} \\
V^b &= \frac{b + \frac{q(\theta)}{\theta}\left[1 - G(-\mu)\right]^{1-\gamma}\left(1 - F(\varepsilon^*)\right)\frac{1}{\chi}}{1 - \beta} \\
V^s &= \frac{s + q(\theta)\left[1 - G(-\mu)\right]^{1-\gamma}\left(1 - F(\varepsilon^*)\right)\frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(\frac{\mu}{\sigma}\right)} + \chi}}{1 - \beta}.
\end{aligned}
$$

Finally optimal pricing, optimal purchase decisions, and optimal entry decisions imply:

$$\begin{aligned}
\varepsilon^* &= b + \beta V^b + p - V^h \\
p &= s + \beta V^s + \frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(\frac{\mu}{\sigma}\right)} + \chi} \\
c^* &= V^h - \left(V^s + (1-L)\,V^b + LV^0\right) \\
k^* &= V^r - V^b.
\end{aligned}$$

One can plug these equations into those for ε^*, c^*, and k^* to get a three equation system with three unknowns that has a unique solution.

For staggered price variant of the model with N staggered groups of list-price-setting sellers, there are a few more variables that require steady state values. Although each group has a different price, in steady state $p^\tau = p$, $\varepsilon^\tau = \varepsilon$, and $V^{s,\tau} = V^s \ \forall \ \tau$, where variables τ are the frictionless steady-state values. The only difference between the two models' steady states, then, is the S^τ. One can show that:

$$S^\tau = \frac{\left[1 - q\left(\theta\right)\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right)\right]^\tau}{\sum_{j=0}^{N-1}\left[1 - q\left(\theta\right)\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right)\right]^j} S,$$

where S is the steady state mass of sellers for the frictionless model. Since all vintages of sellers post the same price in steady state, this does not affect the three equation system for the frictionless model, which continues to define the steady state for the staggered-pricing model.

For the backward-looking variant of the model, the steady state of the model is identical to the frictionless case if $p^i = p$ and $\varepsilon^i = \varepsilon$, where $i = \{N, R\}$.

A.4.7 Simulation Details

I present two main types of results: impulse response functions and stochastic simulations. For both, I report price, sales volume, months of supply, and buyer and seller entry. This section describes how these are computed.

Price is the average transaction price in the market. This is a weighted average of prices, where the weights are equal to the share of transactions accounted for by sellers with each price:

$$p_t = \frac{\int p^i d\left(p_t^i, \Omega_t, \tilde{\theta}_t\right) d\Omega_t}{\int d\left(p_t^i, \Omega_t, \tilde{\theta}_t\right) d\Omega_t}.$$

For a frictionless model, this is just the common price. For the backward-looking model, this is equal to

$p_t = \frac{\alpha d\left(p_t^N,\Omega_t,\tilde{\theta}_t\right)p^N+(1-\alpha)d\left(p_t^R,\Omega_t,\tilde{\theta}_t\right)p_t^R}{\alpha d\left(p_t^N,\Omega_t,\tilde{\theta}_t\right)+(1-\alpha)d\left(p_t^R,\Omega_t,\tilde{\theta}_t\right)}$. For the staggered model, this is equal to

$p_t = \frac{\frac{S_t^0}{S_t^0+S_t^1}d\left(p_t^0,\Omega_t,\tilde{\theta}_t\right)p_t^0+\frac{S_t^1}{S_t^0+S_t^1}d\left(p_t^1,\Omega_t,\tilde{\theta}_t\right)p_t^0}{\frac{S_t^0}{S_t^0+S_t^1}d\left(p_t^0,\Omega_t,\tilde{\theta}_t\right)+\frac{S_t^1}{S_t^0+S_t^1}d\left(p_t^1,\Omega_t,\tilde{\theta}_t\right)}$.

Sales volume is the total amount of sales. This is equal to:

$$vol_t = E_{\Omega_t}\left[d\left(p_t^\tau,\Omega_t,\tilde{\theta}_t\right)\right]S_t,$$

and months of supply is equal to the mass ratio of the mass of sellers to sales volume:

$$MS_t = \frac{vol_t}{S_t} = E_{\Omega_t}\left[d\left(p_t^\tau,\Omega_t,\tilde{\theta}_t\right)\right].$$

Buyer and seller entry come from the laws of motion (48) and (50) and are equal to:

$$\begin{aligned} BEnter_t &= \lambda^r\left(1-K\left(k_t^*\right)\right)R_t + \left(1-LK\left(k_t^*\right)\right)\lambda^h\left(1-C\left(c_t^*\right)\right)H_t \\ SEnter_t &= \lambda^h\left(1-C\left(c_t^*\right)\right)H_t. \end{aligned}$$

Buyer entry has two terms reflecting the entry decisions of renters who receive shocks and new entrants to the metropolitan area.

For the impulse response functions, I use Dynare to compute the impulse response as the average difference between two sets of 100 simulations that use the same sequence of random shocks except for one period in which an additional standard deviation shock is added. For the entry simulation in Figure 1.6.2, I plot the above values at monthly frequencies form a selected simulation. For Table 8, I run 200 simulations of 500 years, collapse the data to quarterly frequency, then calculate the annual change in log price, log volume, and log inventory or estimate the price change on inventory levels regression for each of the 200 simulated series. I report the average values in the table.

A.4.8 Understanding Momentum Arising From Staggered Prices

The full dynamic intuition with staggered pricing is more nuanced than the static intuition presented in the main text because the seller has to weigh the costs and benefits of perturbing price across multiple periods. The intuition is clearest when one considers why a seller does not find it optimal to deviate from a slowly-adjusting price path by listing his or her house at a level closer to the new long-run price after a one-time permanent shock to fundamentals.

After a positive shock to prices, if prices are rising slowly why do sellers not list at a high price, sell at that high price in the off chance that a buyer really likes their house, and otherwise wait until prices are

higher? Search is costly, so sellers do not want to set a very high price and sit on the market for a very long time. Over a shorter time horizon, the probability of sale and profit are very sensitive to perturbing price when a house's price is relatively high but relatively insensitive to perturbing price when a house's price is relatively low. This is the case for two reasons. First, despite the fact that the probability of sale is lower when a house's price is relatively high, demand is much more elastic and so a seller weights that period's low optimal price more heavily. Second, on the equilibrium path, prices converge to steady state at a decreasing rate, so the sellers lose more buyers today by setting a high price than they gain when they have a relatively low price tomorrow. Consequently, in a rising market sellers care about not having too high of a price when their price is high and do not deviate by raising prices when others are stuck at lower prices.

After a negative shock to prices, if prices are falling slowly and search is costly, why do sellers not deviate and cut their price today to raise their probability of sale and avoid search costs if selling tomorrow means selling at a lower price? Although the fact that the elasticity of demand is higher when relative price is higher makes the seller care more about not having too high of a relative price when their price is higher, there is a stronger countervailing effect. Because prices converge to steady state at a decreasing rate on the equilibrium path, sellers setting their price today will undercut sellers with fixed prices more than the sellers are undercut in the future. They thus gain relatively fewer buyers by having a low price when their price is relatively high and leave a considerable amount of money on the table by having a low price when their price is relatively low. On net, sellers care about not having too low of a price when they have the lower price and do not deviate from a path with slowly falling prices.

These intuitions in Figure 38, which shows simulation results for the $N = 2$ case that depict the weights placed the optimal flexible price in each period, the optimal flexible price in each period, and the reset price, equal to the weighted average of the optimal flexible prices. The solid blue line corresponds to the period in which the price is reset, in which the seller is in the higher priced group, and the dashed red line corresponds to the period after the price is reset, in which the seller is in the lower priced group. Panels A, B, and C are for an upward shock, while panels D, E, and F are for a downward shock. In Panel A, the weight is much larger on the reset period—and a lower optimal reset price—because the elasticity of demand is higher. The opposite is true in panel D, as the weighting effect works against momentum, which is why momentum in reset prices appears stronger on the upside in the staggered variant of the model. However, the weighting effect is more made up for by the speed of convergence to steady state, which can most clearly be seen in the asymmetry of the optimal reset price responses in panel E.

Another way of putting these intuitions is that the model features a trade-off between leaving money on the table when a seller has the relatively low price and gaining more buyers when a seller has the relatively high price. On the upside, since price resetters raise prices more than future price setters and since they care

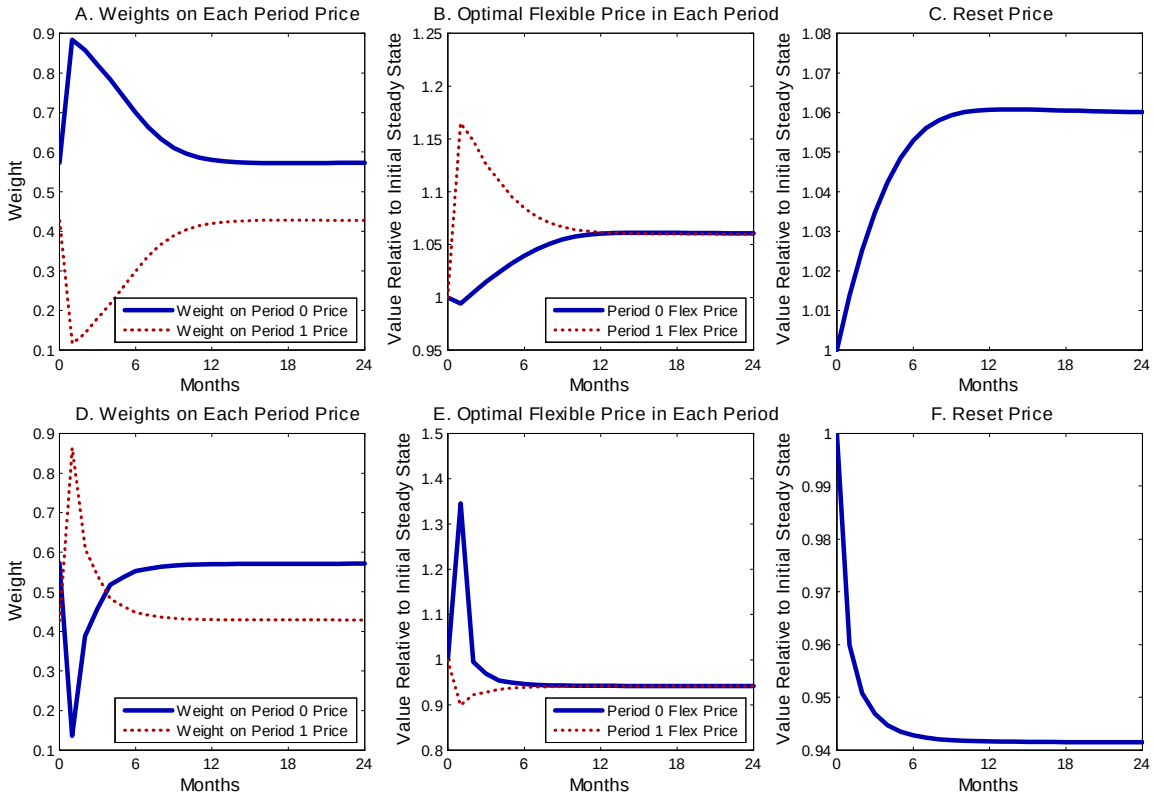

Figure 38: Understanding Momentum Arising From Staggered Pricing

Notes: Panels A, B and C show the impulse response to an upward shock and panels D, E, and F show the impulse response to a downward shock in response to a one-time surprise shock at time zero in a deterministic model solved in Dynare by Newton's method. Panels C and F show the optimal reset prices. Because the optimal reset price is a weighted average of the optimal flexible prices that would prevail in the period the price is reset (period 0) and the period after (period 1), panels A and D show the weights and panels B and E show the optimal flexible prices for each period. Panels C and F are thus equal to the sum of the products of the blue lines and dotted red lines in the preceding two panels.

more about states with more elastic demand, the loss from losing buyers when a resetters have the relatively high price is stronger. On the downside, since price resetters cut prices more than future price resetters, the money left on the table by having a lower price when their prices are relatively low is stronger.

A.5 Calibration

A.5.1 Calibration Targets

The aggregate moments and parameters chosen from other papers are:

- A long-run homeownership rate of 65 percent. The homeownership hovered between 64 percent and 66 percent from the 1970s until the late 1990s before rising in the boom of the 2000s and falling

afterwards.

- $\gamma = 0.8$ from the median specification of Genesove and Han (2012). Anenberg and Bayer (2013) find a similar number.
- $L = 0.7$ from the approximate average internal mover share for Los Angeles of 0.3 from Anenberg and Bayer (2013), which is also roughly consistent with Wheaton and Lee's (2009) analysis of the American Housing Survey and Table 3-10 of the American Housing survey, which shows that under half of owners rented their previous housing unit.
- A median tenure for owner occupants of approximately nine years from American Housing Survey 1997 to 2005 (Table 3-9).
- The approximately equal time for buyers and sellers is from National Association of Realtors surveys (Head et al., 2014; Genesove and Han, 2012). This implies that a normal market is defined by a buyer to seller ratio of $\theta = 1$. I assume a time to sale in a normal market of four months for both buyers and sellers. There is no definitive number for the time to sale, and in the literature it is calibrated between 2.5 and six months. The lower numbers are usually based on real estate agent surveys (*e.g.*, Genesove and Han, 2012), which have low response rates and are effectively marketing tools for real estate agents. The higher numbers are calibrated to match aggregate moments (Piazzesi and Schneider, 2006). I choose four months, which is slightly higher than the realtor surveys but approximately average for the literature.
- Price is equal to \$750,000. The average log price in the IV sample corresponds to a price of \$632,000, so the \$750,000 number implies that on average from 2008 to 2013 prices were 16 percent below their steady-state level. This is on the conservative side for coastal California where the bust was severe and appears to have overshot the long-run equilibrium. The results are not sensitive to using a slightly larger number.
- λ^h and λ^r are set so renters and homeowners experience a shock every 29 months. According to the American Housing Survey 1997 to 2005 (Table 4-9), approximately 41 percent of renters moved within the last year. That translates to a 29 month interval between moves, which I assume are the shocks that induce renters to consider owning. For homeowners, there is not similar data on considered moves, so I assume the same shock probability as renters. This means that homeowners stay in their homes when they receive a shock with higher probability than renters.
- c^*, the amount that the marginal homeowner in steady state would pay to avoid moving and stay in their current house, is set equal to the average transaction cost of selling a home from Haurin and Gill (2002). Haurin and Gill use variation in the length of time that one will stay in a location

among members of the military for whom assignments to a base are known in advance to estimate moving costs of owner-occupancy in a user cost framework. Haurin and Gill's preferred number is three percent of the home's value and four percent of household earnings. However, I do not have household earnings so I use their alternate estimate of five percent of the home's value. I apply this to the steady state price, $c^* = 0.05 \times 750,000 = \$37,500$.

- k^*, the amount the marginal renter in steady state would pay to avoid moving and stay a renter, is set equal to -$20,000 from the imputed tax savings of owner occupancy from Poterba and Sinai (2008). They find that the average household with $125,000 to $250,000 in annual income (I choose this group because their houses are closest in value to average home in my sample) had an annual 2003 tax savings of $7,689 from homeownership ($2,703 from the mortgage interest deduction, $1,125 from property tax deduction, and $3,861 from exclusion of imputed rental income). The average renter in my sample rents for 29 months until the next time they decide to buy or rent, so the total capitalized value of tax savings is just over $18,500 in 2003 dollars, which I adjust to approximately $20,000 in 2008-2013.
- A seven percent annual discount rate corresponding to the mean estimated value for housing searchers in Carrillo (2012). The results are not substantially changed by using values between four and 10 percent.

There are several parameters I set to reasonable values that do not have an important effect on the dynamics:

- The probability a homeowner purchases a home they inspect in steady state is 0.5. This pins down ξ. This parameter does not affect the results unimportant and is set so that the probability of a match is on $[0, 1]$ in the stochastic simulations (with the exception of a few extremely rare circumstances).
- h is set so that the present discounted value of the flow utility of living in a home is approximately 2/3 of its value in steady state, which implies $h = \$7,500$ per month for a $750,000 house. This parameter is a normalization.

Three time series moments are used:

- The persistence of the shock $\rho = 0.99$ is chosen to match the persistence of local income shocks from Glaeser et al. (2013). They estimate an ARMA(1,1) process at the city level net of a city fixed effect and linear drift term to back out the shocks that drive housing cycles. Using BEA income data and Home Mortgage Disclosure Act data on the income of actual home buyers, they find an annual persistence of 0.89, which implies a monthly persistence of 0.99. Their estimated moving average coefficient is small.

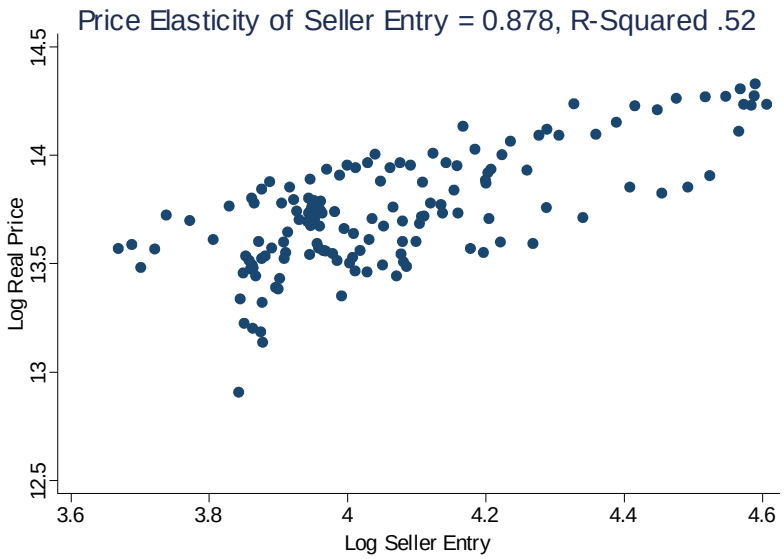

Figure 39: Seller Entry Elasticity

Notes: Seller entry is computed as $\text{Entrants}_t = \text{Sellers}_t - \text{Sellers}_{t-1} + \text{Sales}_t$, with the stock sellers from the Census Vacancy Survey and sales from the National Association of Realtors. Price is the CoreLogic national house price index adjusted for inflation using the CPI and adjusted for seasonality. All data is quarterly from 1976-2013. The regression coefficient of 0.878 is from a regression of log real price on log seller entry.

- A standard deviation of annual log price changes of 0.065 for the real CoreLogic national house price index from 1976 to 2013. This is set to match the standard deviation of aggregate prices for homes that transact collapsed to the quarterly level in stochastic simulations.
- A price elasticity of seller entry of 0.878 based on the CoreLogic, Census, and National Association of Realtors data from 1976 to 2013. As shown in Figure 39, this relationship is very strong in the data. As mentioned in footnote 42 in the main text, I use this moment as a target for both $\bar{c}-\underline{c}$ and $\bar{k}-\underline{k}$ because the stock of buyers is not observed.

Finally, in the calibration to the micro estimates, I use a target of $10,000 for per-period seller flow cost to determine the zero point, which is not precisely estimated in the data and depends on the deadline for a sale used. Together with the target value for the price and the discount rate, the target value for the seller flow cost determines the seller's markup in steady state, which in turn pins down the location of the average price on the demand curve.

The $10,000 target is based on two pieces of evidence from Genesove and Mayer (1997) and Levitt and Syverson (2008):

- Genesove and Mayer show that in their data for Boston condominiums, homes with low equity are listed four percent higher (on a base of $200,000) and are on the market for 15 percent longer at a

100 percent seller loan-to-value ratio relative to an 80 percent seller loan-to-value ratio. They do not report average time on the market, so I assume my average steady-state time on the market of four months. This implies a one month delay nets the seller $13,333.

- Levitt and Syverson show that realtors in Chicago sell their house for on average $7,600 more and remain on the market for 9.5 days longer. This implies that a one month delay nets the seller $23,000, although some of this may be due to harder work on the part of the realtor.[124]

I average these two data points and account for discounting to find that the average seller gives up approximately $18,000 to sell one month sooner. Given that the flow utility of being a homeowner is set to $7,500, I set the flow utility of being a seller to an even -$10,000. This translates in my calibration into a seller markup of $56,500, or roughly 7.5% of the steady-state sales price.

The assumed seller search cost may have an important effect on the degree of momentum in the model because it controls the seller's degree of impatience. To assess the robustness of the results to this parameter, in Appendix A.6.4 I present results from a calibration that uses a much smaller assumed seller search cost of $1,450. Although there is somewhat less momentum because sellers are more patient and willing to forgo matching with a buyer today to obtain a higher price in the future, the model still generates significant momentum.

The 0.4 figure for the AR(1) coefficient in the backward-looking model is described in the main text.

A.5.2 Estimation and Calibration Procedure

As described in the text, the estimation and calibration procedure proceeds in two steps. First, I calibrate to the micro estimates. Then I match the aggregate and time series moments.

Calibration To Micro Estimates

The procedure to calibrate to the micro estimates is largely described in the main text. I start with the IV binned scatter plot (p_b, d_b), which can be thought of as an approximation of the demand curve by 25 indicator functions after the top and bottom 2.5 percent of the price distribution is dropped. In Figure 5, the log relative markup p_b is in log deviations from the average, and I convert it to a dollar amount using the average log price in the IV sample. For each combination of σ, χ, and μ, I use equation (25) to calculate the sum of squared errors:

$$\Sigma_b w_b \left(d_b - d^{3\ month}\left(p_b\right)\right)^2.$$

[124] My estimates imply similar numbers—a one percent price increase on a base of $750,000 increases time on the market by about six to eleven days. My model, however, implies that a one month delay should require a slightly smaller price change than three times a 10 day delay.

Because the data is in terms of probability of sale within 13 weeks, $d^{3\ month}(p_b) = d(p_b) + (1 - d(p_b))\, d(p_b) + (1 - d(p_b))^2\, d(p_b)$ is the simulated probability a house sells within three months. For the weights w_q, I use a normal kernel centered at the average p_b with standard deviation equal to the standard deviation of p_b to ensure I do not over-fit the outliers and consequently over-fit the curvature in the data. I also need to set κ_t, the multiplicative constant. I do so by minimizing the same sum of squared errors for a given vector of the parameters (σ, χ, μ). I then search over the parameter vector to find the (σ, χ, μ) that minimizes the sum of squared errors, using 100 different random starting values.

As mentioned in the main text, I also choose the price that corresponds to the average price in the market $E[p]$ to match a seller search cost of \$10,000 per month. I do so because the zero point in the binned scatter plot is not precisely estimated. The seller surplus pins down the steady-state markup because in steady state:

$$p = \frac{s}{1-\beta} + Markup\left(\frac{\beta}{1-\beta}\Pr[Sell] + 1\right),$$

so given the assumed steady-state values for p and $\Pr[sell]$ a given s pins down the markup. I then find the relative price corresponding to $E[p]$ so that the above calibration procedure matches the implied steady-state markup of \$56,500.

Matching the Aggregate Targets

To match the aggregate targets in Table 6, I invert the steady state so that 16 parameters can be solved for in terms of 16 targets conditional on (σ, χ, μ). I solve this system, defined below, conditional on the 14 steady-state targets and values for the final two target, $\bar{c}-\underline{c}$ and $\bar{k}-\underline{k}$, which are set equal as described above. I then select a value for the standard deviation of innovations to the AR(1) shock σ_η, run 25 random simulations on 500 years of data, and calculate the standard deviation of annual log price changes and the entry elasticity in the simulated data. I adjust the target values for $\bar{c}-\underline{c}$ and σ_η and recalibrate the remainder of the moments until I match the two time series moments. For the backward-looking model, I repeat this procedure altering α until the impulse response to the renter flow utility shock peaks after 36 months.

The calibration procedure is repeated separately for the staggered price and backward-looking variants of the model. Although the results are similar, the σ_η needed to match the data is slightly larger in the backward-looking model due to the additional momentum. While the calibrated values are not the same, although the differences are minor.

The 16 Equation System

Many variables can be found from just a few target values, and I reduce the 16 unknowns to a four equation and four unknown system. I assume there are Per periods per month for the calibration. The

system is defined by:

- β, L, γ, λ^r, λ^h, and V^0 to their assumed values. Note that β, λ^r, and λ^r are adjusted accordingly based on the number of periods in a month.
- $\theta = 1$ from the equality of buyer and seller time on the market.
- The buyer purchases 1/2 of the time, which implies $1 - F(\varepsilon^*) = \frac{1}{2}$. Then using the definition of d and $\theta = 1$,

$$\xi = \frac{1}{4Per\left(1 - G\left(-\mu\right)\right)^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right)}.$$

  where $\left[1 - G\left(-\mu\right)\right]^{1-\gamma}$ can be found using the calibrated value of γ and (σ, μ).
- Given ξ, one can solve for ε^* from $\Pr\left[Sell\right] = \xi \exp\left(-\chi\varepsilon^*\right)\left[1 - G\left(-\mu\right)\right]^{1-\gamma}$.
- The homeownership rate in the model, $\frac{H}{H+B+R}$, is matched to the target moment. Plugging in steady-state values gives:

$$\text{Homeownership Rate} = \frac{\lambda^r\left(1 - K\left(k^*\right)\right)q\left(\theta\right)\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right)}{\begin{array}{c}\lambda^r\left(1 - K\left(k^*\right)\right)q\left(\theta\right)\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right) \\ +\theta\lambda^h\left(1 - C\left(c^*\right)\right)\lambda^r\left(1 - K\left(k^*\right)\right) \\ +LK\left(k^*\right)\lambda^h\left(1 - C\left(c^*\right)\right)q\left(\theta\right)\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right)\end{array}}.$$

  The exogenous target value for the probability a homeowner moves in steady state $\lambda^h\left(1 - C\left(c^*\right)\right)$ (which depends on Per), the known value for $q\left(\theta\right)\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right)$ from time to sale in steady state (which depends on Per), and the target value for λ^r can be used to solve for the value of $K^*\left(k\right)$ that matches the target homeownership rate $Rate$:

$$K\left(k^*\right) = \frac{1}{1 + \frac{1}{\lambda^r}\frac{L \times Rate\lambda^h(1-C(c^*))q(\theta)[1-G(-\mu)]^{1-\gamma}(1-F(\varepsilon^*))}{(1-Rate)q(\theta)[1-G(-\mu)]^{1-\gamma}(1-F(\varepsilon^*))-Rate\times\theta\lambda^h(1-C(c^*))}}.$$

- The population Pop can then be solved for from $Pop = H + B + R$:

$$\begin{aligned} Pop &= \frac{q\left(\theta\right)\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right)}{q\left(\theta\right)\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right) + \lambda^h\left(1 - C\left(c^*\right)\right)} \\ &\quad \times\left(1 + \frac{LK\left(k^*\right)\lambda^h\left(1 - C\left(c^*\right)\right)}{\lambda^r\left(1 - K\left(k^*\right)\right)} + \frac{\lambda^h\left(1 - C\left(c^*\right)\right)}{\frac{q(\theta)}{\theta}\left[1 - G\left(-\mu\right)\right]^{1-\gamma}\left(1 - F\left(\varepsilon^*\right)\right)}\right). \end{aligned}$$

  Again many of the steady-state probabilities depend on Per.

- $(1 - C(c^*))$ can be solved for from the assumed value for λ^h and the probability a homeowner moves in steady state $\lambda^h (1 - C(c^*))$.
- Given targets for $K(k^*)$ and $C(c^*)$ and target values for c^* and k^* as well as $\bar{c}-\underline{c}$ and $\bar{k}-\underline{k}$, the properties of the uniform distribution can be used to back out $\bar{c}$, $\underline{c}$, $\bar{k}$, and $\underline{k}$

$$\begin{aligned} \underline{c} &= c^* - C(c^*)(\bar{c} - \underline{c}) \\ \underline{k} &= k^* - K(k^*)\left(\bar{k} - \underline{k}\right). \end{aligned}$$

This leaves h, u, b, and s, which are solved for jointly to match the target price and satisfy three equilibrium conditions for steady state:

$$\begin{aligned} \varepsilon^* &= b + \beta V^b + p - V^h \\ p &= s + \beta V^s + \frac{1}{\frac{1}{\sigma}\frac{1}{1+\exp\left(\frac{\mu}{\sigma}\right)} + \chi} \\ c^* &= V^h - \left(V^s + (1-L)V^b + LV^0\right) \\ k^* &= V^r - V^b. \end{aligned}$$

A.5.3 Additional Calibrated Values

The calibrated values for the backward-looking model are in Table 7 in the main text. The values for the staggered model are listed in Table 36.

Table 36: Calibrated Parameter Values for Staggered Price Model

Parameter	Value	Parameter	Value	Parameter	Value
β	0.994	$\bar{k}$	\$407k	b	-\$92.4k
γ	0.800	$\underline{k}$	-\$1,160k	s	-\$9.8k
ξ	0.506	Pop	1.484	χ	.0023
λ^h	0.035	L	0.700	σ	3.80
λ^r	0.035	V^0	\$2,606k	μ	10.47
$\bar{c}$	\$458k	h	\$7.5k	σ_η	0.311
$\underline{c}$	-\$1,109k	u	\$3.6k	ρ	0.990

A.6 Additional Simulation Results

A.6.1 Non-Concave Benchmarks

In the main text, I report that without concavity, it would take between 78 and 93 percent backward-looking sellers to generate a three year impulse response to the renter flow utility shock. To generate these numbers, I set $\mu = 300$ (in thousands of dollars) so that the concavity kicks in so far from any equilibrium that the demand curve is linear.

I consider two polar cases. In Figure 10 in the main text, I assume that χ—which is now the constant semi-elasticity of demand since there is no concavity—is equal to its value at the average steady-state price with concavity, $\frac{1}{56.5}$. Under this calibration, 78 percent backward-looking sellers are needed for a 36-month impulse response to the renter flow utility shock. This calibration assumes that the variance of idiosyncratic preference is much smaller than I estimate it to be in my model based on the semi-elasticity of demand for relatively low-priced homes. I use this calibration for the non-concave 26.5 percent backward-looking lines in Figures 10 and 40.

The opposite case I consider is to assume that χ is the same as I estimate in the data. This calibration implies seller markups are much higher but the distribution of buyer idiosyncratic preference is unchanged. Under this calibration, 93 percent backward-looking sellers are needed for a 36-month impulse response to the renter flow utility shock.

Both of these numbers are under 100 percent because with all backward-looking sellers, the impulse response increases without bound. As the fraction of backward-looking sellers is reduced from 100 percent, mean reversion is introduced into the IRF. Consequently, any amount of momentum can be obtained by reducing the fraction of backward-looking sellers from 100 percent.

A.6.2 Downward Shock

Figure 40 shows the impulse response to a downward shock directly analogous to Figure 10. For the 26.5-percent backward-looking model in panel B of Figure 10, and the variants of the staggered model without momentum in panel A, the results look very similar to the upward shock. There is, however, a larger price drop on impact and less momentum for the staggered and concave model. This asymmetry between and upward shock and a downward shock occurs because in the staggered pricing model, groups of sellers leapfrog one another each period, as discussed in Appendix A.4.8. The optimal price is a weighted average of the optimal prices in each period, with the weights corresponding to the discounted probability of sale times the semi-elasticity of demand. The semi-elasticity of demand—and hence the weight on price—is higher in the period with a lower price, which helps momentum for an upward shock but works against momentum and

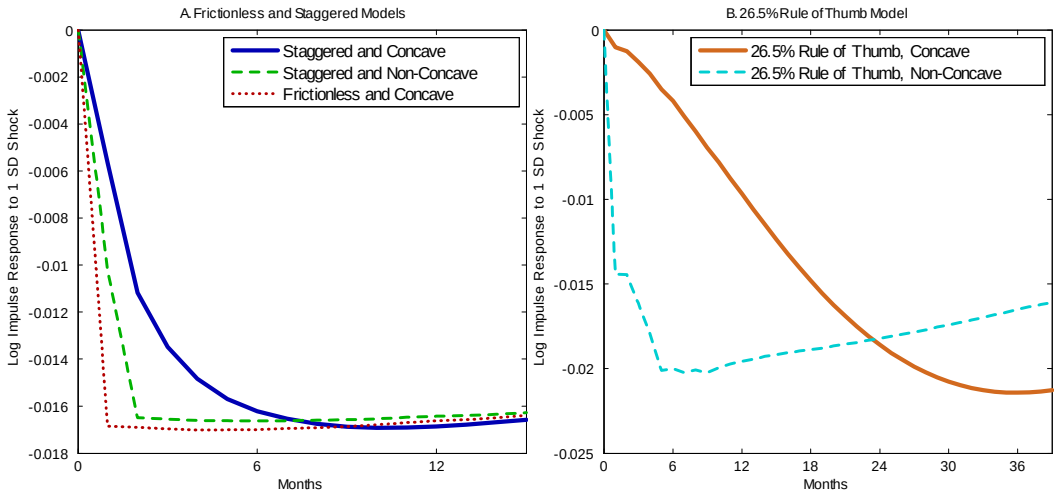

Figure 40: Price Impulse Response Functions: Downward Shock

Notes: Panel A shows the impulse responses to a one standard deviation shock to the flow utility of renting in the frictionless model with concave demand, the staggered model with concave demand, and the staggered model without concave demand. For the model without concave demand, the threshold for being overpriced μ is raised to a level that is never reached, the slope of the demand curve is adjusted to the steady-state slope at the average price in the concave model, the model is recalibrated, and the standard deviation of the stochastic shock is adjusted so that the impulse response is even with the frictionless and concave impulse response after a year. Panel B shows the impulse responses to a one standard deviation shock to the flow utility of renting in the backward-looking model with and without concavity. Simulated impulse responses are calculated by differencing two simulations of the model from periods 100 to 150, both of which use identical random shocks except in period 101 in which a one standard deviation negative draw is added to the random sequence and then computing the average difference over 100 simulations.

causes the larger drop on impact in for a downward shock.

A.6.3 Impulse Responses For All Variables

In the main text, I only show impulse-responses for non-price variables for an upward shock in the 26.5 percent backward-looking model in Figure 11. Figures 42, 43, and 44 show the same impulse responses for a downward shock in the 26.5 percent backward-looking model and for an upward and downward shock in the staggered models.

The downward shock for the backward-looking model looks close to the mirror image of the upward shock. The staggered model, however, looks somewhat different. While there is still a gradual price response and a shorter-lived but analogous buyer and seller entry response, volume and inventory both spike on impact. This is because prices change more rapidly, so buyers in the market when the shock occurs have a strong incentive to transact today for an upward shock or to wait to buy for a downward shock. While months of supply still overshoots, this effect is less substantial and the spike is more prominent.

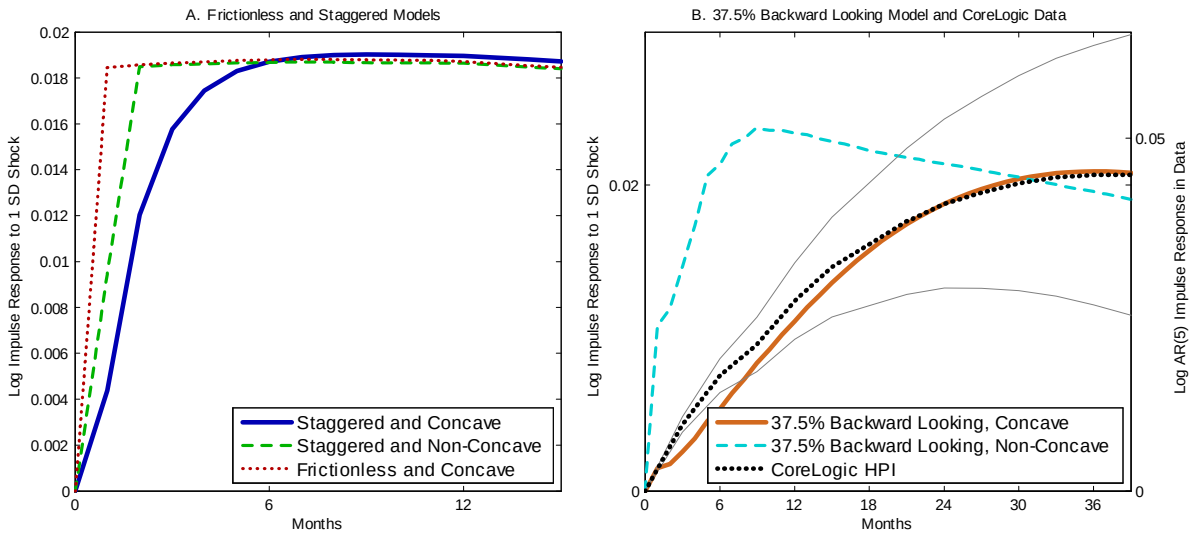

Figure 41: Impulse Responses For Calibration With Low Seller Search Cost

Notes: This Figure is the analogue of Figure 10 in the main text using a lower seller search cost for the calibration. Panel A shows the impulse responses to a one standard deviation negative shock to the flow utility of renting in the frictionless model with concave demand, the staggered model with concave demand, and the staggered model without concave demand. For the model without concave demand, the threshold for being overpriced μ is raised to a level that is never reached, the slope of the demand curve is adjusted to the steady-state slope at the average price in the concave model, the model is recalibrated, and the standard deviation of the stochastic shock is adjusted so that the impulse response is even with the frictionless and concave impulse response after a year. Panel B shows the impulse responses to a one standard deviation shock to the flow utility of renting in the backward-looking model with and without concavity. For the model without concave demand, the threshold for being overpriced μ is raised to a level that is never reached, the slope of the demand curve is adjusted to the steady-state slope at the average price in the concave model, and the model is recalibrated. Also shown in panel B in the dotted black line and with grey 95% confidence intervals and on the right axis is the impulse response to a one standard deviation price shock estimated from a quarterly AR(5) for the seasonally and CPI adjusted CoreLogic national house price index for 1976-2013, as in Figure 1. Simulated impulse responses are calculated by differencing two simulations of the model from periods 100 to 150, both of which use identical random shocks except in period 101 in which a one standard deviation negative draw is added to the random sequence, and then computing the average difference over 100 simulations.

A.6.4 Calibration Robustness: Lower Seller Search Cost

The calibration procedure assumes a seller search cost of \$10,000 per month based on figures from two papers. This assumed parameter is important for the degree of amplification of momentum in the model because it controls the degree to which sellers are willing to forgo a higher price in the future in order to attract buyers today. To assess the robustness of the results to this important assumed parameter, this section presents a calibration with a far smaller seller search cost of \$1,450.

To do so, I use the calibration procedure detailed in Appendix A.5 but change the assumed seller search cost to \$1,450. This results in similar parameter values for most variables except μ, which is smaller as the calibration procedure shifts the average price further into the elastic region of the demand curve.

The results are shown in Figure 41. Two months of staggered pricing leads to an impulse response with seven to eight months of momentum instead of 10 months in the baseline calibration. The variant of the model with backward-looking sellers generates a three-year impulse response with 37.5 percent of backward-looking sellers, rather than 26.5 percent in the baseline calibration. Thus while using a smaller assumed seller search cost does reduce the degree of momentum generated by the model, the calibrated model still generates substantial amplification of the underlying shock with a much smaller seller search cost.

A.6.5 Deterministic Impulse Responses For Staggered Model

To show that my results are not due to error in the log-cubic approximation pruning higher order terms used in the stochastic simulations, Figure 45 shows the impulse response to an upward and downward unexpected one-time deterministic shock solved exactly by Newton's method for the staggered model. These impulse responses look similar to their stochastic counterparts.

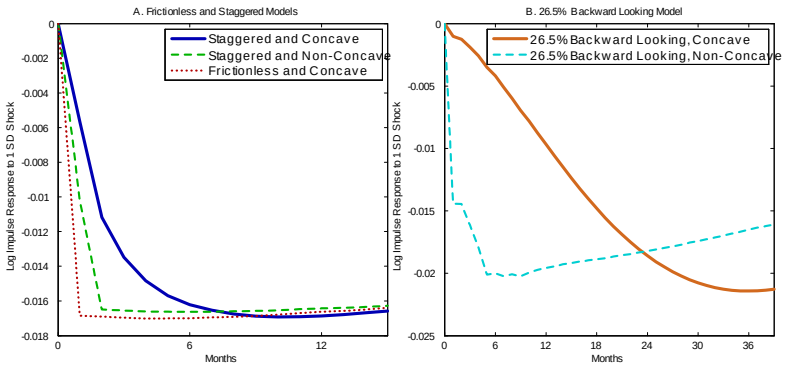

Figure 42: Impulse Responses to Downward Shock in the Rule-of-Thumb Model

Notes: Each panel plots the indicated impulse response to a one standard deviation downward shock for the frictionless and backward-looking variants of the model. The frictionless model uses the same calibration and shock as the 26.5 percent backward-looking model with no backward-looking sellers. Simulated impulse responses are calculated by differencing two simulations of the model from periods 100 to 150, both of which use identical random shocks except in period 101 in which a one standard deviation negative draw is added to the random sequence and then computing the average difference over 100 simulations.

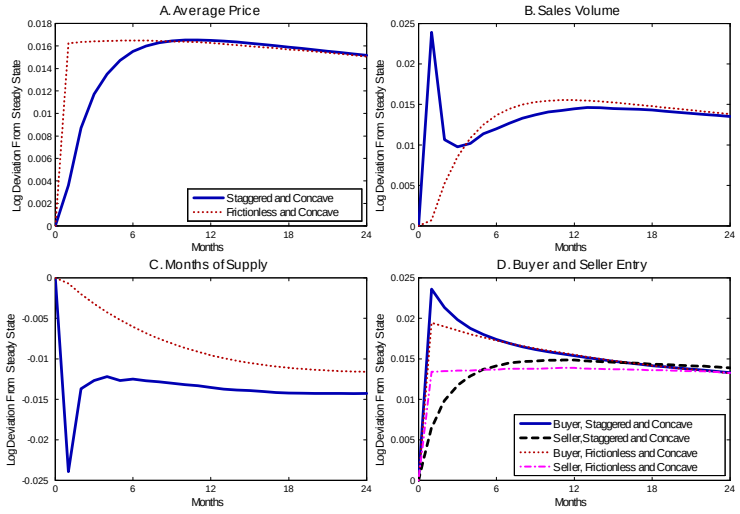

Figure 43: Impulse Responses to Downward Shock in the Rule-of-Thumb Model

Notes: Each panel plots the indicated impulse response to a one standard deviation upward shock for the frictionless and staggered variants of the model. The frictionless model uses the same calibration and shock but all sellers set their price simultaneously. Simulated impulse responses are calculated by differencing two simulations of the model from periods 100 to 150, both of which use identical random shocks except in period 101 in which a one standard deviation negative draw is added to the random sequence and then computing the average difference over 100 simulations.

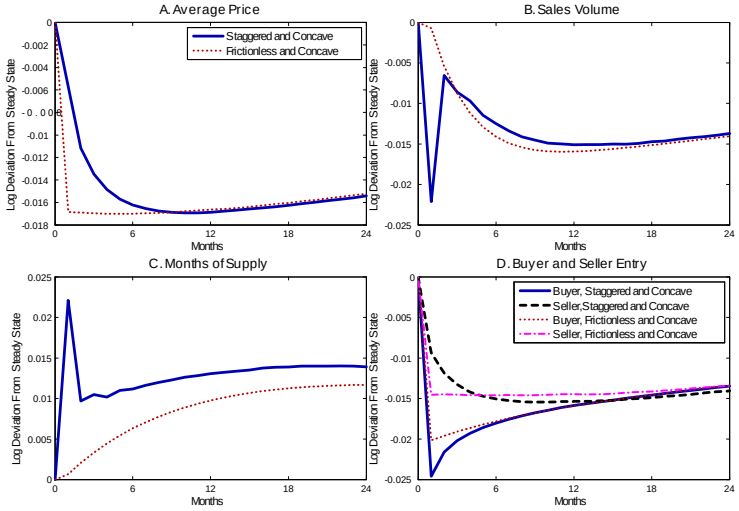

Figure 44: Impulse Responses to Downward Shock in the Rule-of-Thumb Model

Notes: Each panel plots the indicated impulse response to a one standard deviation downward shock for the frictionless and staggered variants of the model. The frictionless model uses the same calibration and shock but all sellers set their price simultaneously. Simulated impulse responses are calculated by differencing two simulations of the model from periods 100 to 150, both of which use identical random shocks except in period 101 in which a one standard deviation negative draw is added to the random sequence and then computing the average difference over 100 simulations.

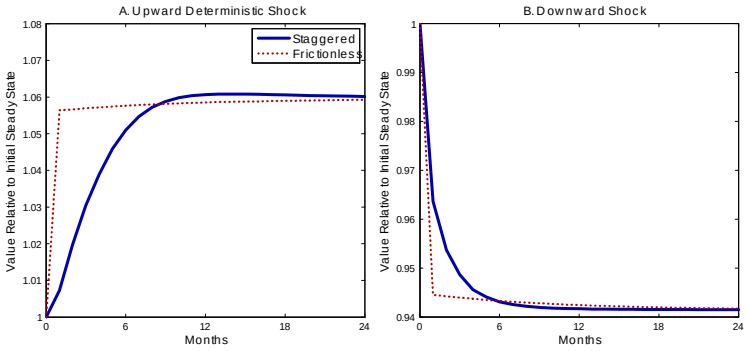

Figure 45: Response to One-Time Deterministic Shock (Non-Approximated)

Notes: Each panel plots the value of the indicated variable divided by its initial steady-state value in response to a -.3 deterministic permanent shock to u (roughly consistent with one standard deviation stochastic shock) in the upward panel and a .3 deterministic shock in the downward panel. The frictionless model uses the same calibration and shock but all sellers set their price simultaneously. Both are solved exactly by Newton's method assuming that the model returns to steady state in 500 years and use the same calibration as the stochastic staggered model, and the frictionless model includes no staggering.

B Chapter 2 Appendix

B.1 Derivations and Proofs

B.1.1 Steady Sate

Denoting steady state values with a star superscript, the no default assumption means that $l_1^* = 0$, $v_d^* = 0$, $v_r^* = 0$. $v_n^* = v_b^*$ and $l_0^* = 1 - v_n^*$. This implies $\mu^* = 1$, $q_s(\mu^*) = q_b(\mu^*) = \chi$, $r_n^* = 1$, and $r_d^* = 0$. Because inflows into being a seller and outflows from being a seller are equal in steady state,

$$\begin{aligned}(1 - v_n^*)\gamma &= v_n^* \chi (1 - F(h_n^*)) \\ v_n^* &= \frac{\gamma}{\gamma + \chi(1 - F(h_n^*))} \quad l_0 = \frac{\chi(1 - F(h_n^*))}{\gamma + \chi(1 - F(h_n^*))}.\end{aligned}$$

Replacing the conditional expectations of surpluses with differences of cutoffs as in (38) and setting $r_n^* = 1$, $r_d^* = 0$, and $q_s^* = q_b^* = \chi$ yields simplified steady state value functions:

$$\begin{aligned}V_h^* &= \frac{h + \beta\{\gamma(V_n^* + B^*)\}}{1 - \beta(1 - \gamma)} \\ B^* &= \frac{u_b}{1 - \beta} + \frac{(1 - \theta)\chi(1 - F(h_n^*))E[h - h_n^* | h \geq h_n^*]}{(1 - \beta)(1 - \beta(1 - \gamma))} \\ V_m^* &= \frac{m}{1 - \beta} + \frac{\theta\chi(1 - F(h_m^*))E[h - h_m^* | h \geq h_m^*]}{(1 - \beta)(1 - \beta(1 - \gamma))}, m \in \{n, d\} \\ R^* &= \frac{u_r + \beta\sigma B^*}{1 - \beta(1 - \sigma)}.\end{aligned}$$

With everything in terms of the cutoffs, a two-equation system that pins down h_n^* and h_d^*. Subtracting the cutoff condition (34) at the distressed and non-distressed cutoffs gives:

$$V_{h_n}^* - V_{h_d}^* = (m_n - m_d) + \beta[V_n^* - V_d^*].$$

Plugging in the steady state values and manipulating yields an equation that implicitly defines the difference of the cutoffs:

$$\begin{aligned}h_n^* - h_d^* &= (m_n - m_d)\frac{1 - \beta(1 - \gamma)}{(1 - \beta)} \\ &\quad + \theta\frac{\beta}{1 - \beta}\chi\{(1 - F(h_n^*))E[h - h_n^* | h \geq h_n^*] - (1 - F(h_d^*))E[h - h_d^* | h \geq h_d^*]\}.\end{aligned} \tag{57}$$

The second equation comes from evaluating the cutoff condition (34) at h_n^*:

$$V_{h_n^*} = m_n + u_b + \beta[B^* + V_n^*]. \tag{58}$$

Equations (57) and (58) define a system that can be solved for h_d^* and h_n^*. All of the other steady-state variables are written in terms of these cutoffs.

Proposition 8 *If $a < m_n + u_b + \frac{\beta(1-\gamma)\chi E[h-a]}{1-\beta(1-\gamma)}$, there exists a unique steady state of the model.*

Proof. *Because there are no REO sellers in steady state,*

$$V_n^* + B^* = \frac{m_n + u_b}{1-\beta} + \frac{\chi\left(1 - F\left(h_m^*\right)\right) E\left[h - h_m^* | h \geq h_m^*\right]}{\left(1-\beta\right)\left(1-\beta\left(1-\gamma\right)\right)}$$

The cutoff condition for h_n^* is:

$$\begin{aligned}
\frac{h_n^* + \beta\left[V_n^* + B^*\right]}{1-\beta\left(1-\gamma\right)} &= m_n + u_b + \beta\left[V_n^* + B^*\right] \\
\frac{h_n^*}{1-\beta\left(1-\gamma\right)} &= m_n + u_b + \beta\left(\frac{1-\beta\left(1-\gamma\right)-\gamma}{1-\beta\left(1-\gamma\right)}\right)\left[V_n^* + B^*\right] \\
h_n^* &= \left(1-\beta\left(1-\gamma\right)\right)\left(m_n + u_b\right) + \beta\left(1-\beta\right)\left(1-\gamma\right)\left[V_n^* + B^*\right]
\end{aligned}$$

Plugging in for $V_n^* + B^*$ and re-arranging yields:

$$\begin{aligned}
h_n^* &= \left(1-\beta\left(1-\gamma\right)\right)\left(m_n + u_b\right) + \beta\left(1-\beta\right)\left(1-\gamma\right)\left[\frac{m_n + b}{1-\beta} + \frac{\chi\left(1 - F\left(h_n^*\right)\right) E\left[h - h_n^* | h \geq h_n^*\right]}{\left(1-\beta\right)\left(1-\beta\left(1-\gamma\right)\right)}\right] \\
h_n^* &= m_n + u_b + \frac{\beta\left(1-\gamma\right)\chi\left(1 - F\left(h_n^*\right)\right) E\left[h - h_n^* | h \geq h_n^*\right]}{1-\beta\left(1-\gamma\right)} \qquad (59)
\end{aligned}$$

We want to find a unique solution to this equation on $h_n^* \in [a, \infty)$. As $h_n^* \to \infty$, the RHS of equation (59) approaches $m_n - u_b$. As $h_n^* \to a$, the RHS of equation (59) approaches $m_n + u_b + \frac{\beta(1-\gamma)\chi E[h-a]}{1-\beta(1-\gamma)}$. Thus as long as $a < m_n + u_b + \frac{\beta(1-\gamma)\chi E[h-a]}{1-\beta(1-\gamma)}$, since both the RHS and LHS of equation (59) are continuous in h_n^*, by the intermediate value theorem we know there exists a solution on $[a, \infty)$ to equation (59). Furthermore, we know that the LHS is strictly increasing in h_n^*, while the RHS is strictly decreasing in h_n^* since:

$$\frac{d}{dx}\left(1 - F\left(x\right)\right) E\left[h - x | h \geq x\right] = -\left(1 - F\left(x\right)\right) < 0.$$

This implies that the solution to equation (59) is unique.

Finally note that

$$h_d^* \; - \theta \frac{\beta}{1-\beta} \chi\left(1 - F\left(h_d^*\right)\right) E\left[h - h_d^* | h \geq h_d^*\right]$$

is a monotonically increasing function of h_d^* and thus, given the solution h_n^* to equation (59), there exists a unique solution to equation (57). ■

Note that given our assumptions it is generally the case that $a < m_n + u_b + \frac{\beta(1-\gamma)\chi E[h-a]}{1-\beta(1-\gamma)}$ and thus that there is a unique equilibrium. This is because β is close to 1 and γ is close to 0, so the denominator of the

fraction is very small. Uniqueness would only be a concern with a very low discount factor or high moving probability.

Due to higher holding costs and balance sheet concerns, an REO seller should be more willing to sell the property conditional on being matched with a buyer than a normal seller. We show this is always the case in steady state:

Lemma 9 *For a given h, the probability of sale for a distressed seller is higher than the probability of sale for a non-distressed seller.*

Proof. *Note again that:*

$$\frac{d}{dx}\left(1-F\left(x\right)\right)E\left[h-x|h\geq x\right]=-\left(1-F\left(x\right)\right)<0.$$

Suppose that $h_n^* \leq h_d^*$. Then:

$$\begin{aligned}\left(h_n^*-h_d^*\right) & -\theta\frac{\beta}{1-\beta}\chi\left\{\left(1-F\left(h_n^*\right)\right)E\left[h-h_n^*|h\geq h_n^*\right]-\left(1-F\left(h_d^*\right)\right)E\left[h-h_d^*|h\geq h_d^*\right]\right\}\\ < & \quad 0\\ < & \quad \left(m_n-m_d\right)\frac{1-\beta\left(1-\gamma\right)}{\left(1-\beta\right)}.\end{aligned}$$

which contradicts equation (57). It must therefore be that $h_n^* > h_d^*$, which indicates that distressed sellers are more likely to sell than non-distressed sellers. ■

We use the Nash bargaining condition to back out steady state prices. We have for $m \in \{n, d\}$ and a given h:

$$\frac{\theta}{1-\theta}=\frac{S_m^S}{S_m^B}=\frac{p_m^*\left(h\right)-m-\beta V_m^*}{V_h^*-p_m^*\left(h\right)-\beta B^*}$$

And so

$$\begin{aligned}\theta\left[V_h^*-p_m^*\left(h\right)-\beta B^*\right] &= \left(1-\theta\right)\left[p_m^*\left(h\right)-m-\beta V_m^*\right]\\ p_m^*\left(h\right) &= \theta V_h^*-\beta\theta B^*+\left(1-\theta\right)m+\left(1-\theta\right)\beta V_m^*\\ &= \theta V_h^*+m+\beta V_m^*-\theta\left[m+\beta V_m^*+\beta B^*\right]\\ &= \theta\left[V_h^*-V_{h_m^*}\right]+m+\beta V_m^*\\ &= \frac{\theta\left(h-h_m^*\right)}{1-\beta\left(1-\gamma\right)}+m+\frac{\beta m}{1-\beta}+\frac{\beta\theta\chi\left(1-F\left(h_m^*\right)\right)E\left[h-h_m^*|h\geq h_m^*\right]}{\left(1-\beta\right)\left(1-\beta\left(1-\gamma\right)\right)}\\ &= \frac{\theta\left(h-h_m^*\right)}{1-\beta\left(1-\gamma\right)}+\frac{m}{1-\beta}+\frac{\beta\theta\chi\left(1-F\left(h_m^*\right)\right)E\left[h-h_m^*|h\geq h_m^*\right]}{\left(1-\beta\right)\left(1-\beta\left(1-\gamma\right)\right)} \qquad (60)\end{aligned}$$

It is also always the case that distressed properties sell for less than non-distressed properties:

Proposition 10 *Distressed sales trade at a constant discount, in the sense that $p_n^*(h) - p_d^*(h) = \Delta$ for all $h \geq h_n^* > h_d^*$ for some constant $\Delta > 0$.*

Proof. *Taking the difference of prices we get:*

$$\begin{aligned} p_n^*(h) - p_d^*(h) &= \frac{\theta(h_d^* - h_n^*)}{1-\beta(1-\gamma)} + \frac{m_n - m_d}{1-\beta} \\ &\quad + \frac{\beta\theta\chi\left\{(1-F(h_n^*))\,E[h-h_n^*|h \geq h_n^*] - (1-F(h_d^*))\,E[h-h_d^*|h \geq h_d^*]\right\}}{(1-\beta)(1-\beta(1-\gamma))} \end{aligned}$$

Using equation (57) this simplifies to:

$$\begin{aligned} p_n(h) - p_d(h) &= \frac{\theta(h_d^* - h_n^*)}{1-\beta(1-\gamma)} + \frac{m_n - m_d}{1-\beta} - \frac{m_n - m_d}{1-\beta} + \frac{h_n^* - h_d^*}{1-\beta(1-\gamma)} \\ &= \frac{1-\theta}{1-\beta(1-\gamma)}(h_n^* - h_d^*) \equiv \Delta > 0 \end{aligned}$$

■

B.1.2 Dynamics

This appendix describes how we solve the dynamic model with exogenous defaults presented in section 2.3.

First, the laws of motion simply add inflows and subtract outflows according to figure17:

$$\begin{aligned} l_0(t+1) &= (1-\gamma)\,l_0(t) + v_b(t)\,q_b(\mu(t))\sum_m r_m(t)\,e^{-\lambda(h_m(t)-a)} \\ l_1(t+1) &= (1-\gamma)\,l_1(t) \\ v_n(t+1) &= \gamma l_0(t) + \gamma(1-\alpha)\,l_1(t) + v_n\left[1 - q_s(\mu(t))\,e^{-\lambda(h_n(t)-a)}\right] \\ v_d(t+1) &= \gamma\alpha l_1(t) + v_d\left[1 - q_s(\mu(t))\,e^{-\lambda(h_d(t)-a)}\right] \\ v_b(t+1) &= \gamma l_0(t) + \gamma(1-\alpha)\,l_1(t) + v_r(t)\sigma + v_b(t)\left[1 - q_b(\mu(t))\sum r_m(t)\,e^{-\lambda(h_m(t)-a)}\right] \\ v_r(t+1) &= \gamma\alpha l_1(t) + (1-\sigma)v_r(t) \end{aligned}$$

To generate the dynamic path of V_h, we expand the sum in equation (**??**) and collecting terms:

$$V_h(t) = \frac{h}{1-\beta(1-\gamma)} + \beta\sum_{j=1}^{\infty}\left[(\beta(1-\gamma))^{j-1}\left\{\gamma\left[V_n(t+j) + B(t+j)\right]\right\}\right].$$

The sum in the second term can be written recursively as:

$$\Gamma(t) = \gamma\left[V_n(t) + B(t)\right] + \beta(1-\gamma)\,\Gamma(t+1) \tag{61}$$

so that

$$V_h(t) = \frac{h}{1-\beta(1-\gamma)} + \beta\Gamma(t+1).$$

The entire dynamic system can thus be written recursively. The cutoff rule then simplifies to:

$$\frac{h_m(t)}{1-\beta(1-\gamma)} + \beta\Gamma(t+1) = u_b + \beta B(t+1) + m + \beta V_m(t+1). \tag{62}$$

The full general equilibrium model is made up of 13 endogenous variables and 13 equations. The endogenous variables are the cutoffs h_n and h_d, the masses v_n, v_d, v_b, v_r, l_0, and l_1, and value functions Γ, V_{m_d}, V_{m_r}, B, and R. From these values, all of the other endogenous parameters of the model can be determined. Substituting out the conditional expectation of the surplus using (38), using the exponential distribution, and using the definitions of q_b and q_s from the matching function gives the dynamic system:

$$\begin{aligned}
V_n(t) &= \beta V_n(t+1) + m_n + \chi\mu(t)\frac{\theta}{\lambda[1-\beta(1-\gamma)]}e^{-\lambda(h_n(t)-a)} \\
V_d(t) &= \beta V_d(t+1) + m_d + \chi\mu(t)\frac{\theta}{\lambda[1-\beta(1-\gamma)]}e^{-\lambda(h_d(t)-a)} \\
B(t) &= \beta B(t+1) + u_b + \chi\mu(t)\frac{1-\theta}{\lambda[1-\beta(1-\gamma)]}\sum_m r_m(t)\,e^{-\lambda(h_m(t)-a)} \\
R(t) &= u_r + \beta\{\sigma B(t+1) + (1-\sigma)R(t+1)\} \\
\Gamma(t) &= \gamma[V_n(t) + B(t)] + \beta(1-\gamma)\Gamma(t+1) \\
l_0(t+1) &= (1-\gamma)l_0(t) + v_b(t)\chi\mu(t)\sum_m r_m(t)\,e^{-\lambda(h_m(t)-a)} \\
l_1(t+1) &= (1-\gamma)l_1(t) \\
v_n(t+1) &= \gamma l_0(t) + \gamma(1-\alpha)l_1(t) + v_n\left[1-\chi\mu(t)\,e^{-\lambda(h_n(t)-a)}\right] \\
v_d(t+1) &= \gamma\alpha l_1(t) + v_d\left[1-\chi\mu(t)\,e^{-\lambda(h_d(t)-a)}\right] \\
v_b(t+1) &= \gamma l_0(t) + \gamma(1-\alpha)l_1(t) + v_r(t)\sigma + v_b(t)\left[1-\chi\mu(t)\sum r_m(t)\,e^{-\lambda(h_m(t)-a)}\right] \\
v_r(t+1) &= \gamma\alpha l_1(t) + (1-\sigma)v_r(t) \\
\frac{h_n(t)}{1-\beta(1-\gamma)} + \beta\Gamma(t+1) &= u_b + \beta B(t+1) + m_n + \beta V_n(t+1) \\
\frac{h_d(t)}{1-\beta(1-\gamma)} + \beta\Gamma(t+1) &= u_b + \beta B(t+1) + m_d + \beta V_d(t+1)
\end{aligned}$$

where $\mu(t) = \frac{v_b(t)}{v_n(t)+v_d(t)}$ and $r_m(t) = \frac{v_m(t)}{v_n(t)+v_d(t)}$. We solve this system using Newton's Method as implemented in DYNARE, which guesses that the model returns to steady state at time T, solves a system of $13T$ equations, and checks that the model is in fact within ε of the steady state at time T. Solving the model with endogenous defaults of section 2.5 is performed similarly although the laws of motion are modified as described in appendix B.1.4.

We then back out prices as in the steady state. Prices are defined by (39). The mean price for a type m seller is then:

$$\bar{p}_m = E_h\left[p_{m,h}\left(t\right)|h_m \geq h_m\right] = \frac{\theta}{\lambda}\frac{1}{1-\beta\left(1-\gamma\right)} + m + \beta V_m\left(t+1\right) \tag{63}$$

and the overall mean price in a price index is:

$$\begin{aligned}\bar{p} &= FVol_N\left[\frac{\theta}{\lambda}\frac{1}{1-\beta\left(1-\gamma\right)} + m_n + \beta V_n\left(t+1\right)\right] \\ &\quad + FVol_D\left[\frac{\theta}{\lambda}\frac{1}{1-\beta\left(1-\gamma\right)} + m_d + \beta V_d\left(t+1\right)\right]\end{aligned} \tag{64}$$

where $FVol_m$ is the fraction of total volume accounted for by type m sellers.

B.1.3 Calibration

As described in the main text, we use five aggregate moments from the housing market prior to the crash to set a, λ, m_n, and m_d. These moments are the average price of a normal home in steady state $\bar{p}_n^*$, the variance of the residual price distribution $\sigma_{p_n^*}^2$, the discount for a distressed sale in terms of mean prices $\frac{\bar{p}_n - \bar{p}_d}{\bar{p}_n}$, and the time on the market for a normal sale T_n^* and a distressed sale T_d^*

Using the expressions for the price and the probability of sale in the main text along with properties of the exponential distribution, these moments are:

$$\begin{aligned}\bar{p}_n^* &= \frac{\theta}{\lambda\left[1-\beta\left(1-\gamma\right)\right]} + \frac{m_n}{1-\beta} + \frac{\beta\theta\chi e^{-\lambda\left(h_n^*-a\right)}}{\lambda\left(1-\beta\right)\left(1-\beta\left(1-\gamma\right)\right)} \\ \sigma_{p_n^*}^2 &= \frac{\theta^2}{\lambda^2\left[1-\beta\left(1-\gamma\right)\right]^2} \\ \frac{\bar{p}_n - \bar{p}_d}{\bar{p}_n} &= \frac{m_n - m_d}{\bar{p}_n\left(1-\beta\right)} + \frac{\beta\theta\chi\left\{e^{-\lambda\left(h_m^*-a\right)} - e^{-\lambda\left(h_m^*-a\right)}\right\}}{\lambda\bar{p}_n\left(1-\beta\right)\left(1-\beta\left(1-\gamma\right)\right)} = \frac{h_n^* - h_d^*}{\bar{p}_n\left[1-\beta\left(1-\gamma\right)\right]} \\ T_n^* &= \frac{1-\chi\exp\left(-\lambda\left(h_n^*-a\right)\right)}{\chi\exp\left(-\lambda\left(h_n^*-a\right)\right)} = \frac{1}{\chi}\exp\left(\lambda\left(h_n^*-a\right)\right) - 1 \\ T_d^* &= \frac{1-\chi\exp\left(-\lambda\left(h_d^*-a\right)\right)}{\chi\exp\left(-\lambda\left(h_d^*-a\right)\right)} = \frac{1}{\chi}\exp\left(\lambda\left(h_d^*-a\right)\right) - 1\end{aligned}$$

Plugging the second and fourth equations into the first gives:

$$\bar{p}_n^* = \sigma_{p_n^*} + \frac{m_n}{1-\beta} + \frac{\beta}{1-\beta}\frac{\sigma_{p_n^*}}{T_n^*+1}$$

which implicitly defines m_n as a function of known parameters and observable moments.

We then define a six equation system with six variables – a, θ, λ, m_d, h_n^*, and h_d^* – that we use to calibrate the remainder of the model. Taking the square root of the second equation and rearranging gives

λ as a function of θ and observable moments:

$$\lambda = \frac{\theta}{\sigma_{p_n^*}\left[1-\beta\left(1-\gamma\right)\right]}.$$

An expression for a is obtained by inverting the fourth equation $h_n^* = a + \frac{1}{\lambda}\ln\left(\chi\left(T_n^*+1\right)\right)$ and then plugging into the cutoff condition for n:

$$\begin{aligned} V_{h_n^*} &= m_n + u_b + \beta\left[B^* + V_{m_d}^*\right] \\ h_n^* &= \left(1-\beta\left(1-\gamma\right)\right)m_n + \beta\left(1-\gamma\right)\left[m_n + \frac{\chi e^{-\lambda(h_n^*-a)}}{\lambda\left(1-\beta\left(1-\gamma\right)\right)}\right] \end{aligned}$$

Plugging in for h_n^* and solving gives:

$$a = \left(1-\beta\left(1-\gamma\right)\right)m_n + \beta\left(1-\gamma\right)\left[m_n + \frac{1}{\left(T_n^*+1\right)\lambda\left(1-\beta\left(1-\gamma\right)\right)}\right] - \frac{1}{\lambda}\ln\left(\chi\left(T_n^*+1\right)\right)$$

The equations for λ and a, the moments for T_n^*, T_d^*, and the discount,

$$\frac{\bar{p}_n - \bar{p}_d}{\bar{p}_n} = \frac{h_n^* - h_d^*}{\bar{p}_n\left[1-\beta\left(1-\gamma\right)\right]} \qquad T_n^* = \frac{1}{\chi}\exp\left(\lambda\left(h_n^*-a\right)\right) - 1 \quad T_d^* = \frac{1}{\chi}\exp\left(\lambda\left(h_d^*-a\right)\right) - 1,$$

along with (57),

$$h_n^* - h_d^* = \left(m_n - m_d\right)\frac{1-\beta\left(1-\gamma\right)}{\left(1-\beta\right)} + \frac{\theta}{\lambda}\frac{\beta}{1-\beta}\chi\left\{e^{-\lambda(h_n^*-a)} - e^{-\lambda(h_d^*-a)}\right\},$$

form the six equation system, which we solve numerically.

Although all of the variables are jointly determined, we have found that θ and the gap between m_d and m_n are principally determined by the gap in time on the market and the REO discount while a and λ are principally determined by the moments of the price distribution.

B.1.4 Extended Model

For the extended model, the housing market is unchanged and so the value functions are unchanged. Only the laws of motion differ, as described by Figure 20. As described in the text, we have two different exogenous shocks. First, we assume that a fraction $\delta\left(t\right)$ of individuals who sell due to taste shocks become

renters instead of buyers and shock $\delta(t)$. This leads to the following laws of motion:

$$\begin{aligned}
l_0(t+1) &= (1-\gamma) l_0(t) + v_b(t) q_b(\mu(t)) \sum r_m(t) (1 - F(h_m(t))) + w(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} \\
l_1(t+1) &= (1-\gamma) l_1(t) \\
w(t+1) &= (1-\gamma_I) w(t) + (\gamma - \gamma_I) l_1(t) (1 - G(V_n(t))) - w(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} \\
f(t+1) &= \gamma_I l_1(t) (1 - G(V_n(t))) + \gamma_I w(t) \\
&\quad + f(t) \left(1 - \frac{1}{\phi f(t) + 1}\right) - f(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} \\
v_n(t+1) &= \gamma l_0(t) + \gamma l_1(t) G(V_n(t)) + (w(t) + f(t)) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} \\
&\quad + v_n [1 - q_s(\mu(t)) (1 - F(h_n(t)))] \\
v_d(t+1) &= \frac{f(t)}{\phi f(t) + 1} + v_d [1 - q_s(\mu(t)) (1 - F(h_d(t)))] \\
v_b(t+1) &= (1-\delta) [\gamma l_0(t) + \gamma l_1(t) G(V_n(t))] + \\
&\quad + v_r(t)\sigma + v_b(t) \left[1 - q_b(\mu(t)) \sum r_m(t) (1 - F(h_m(t)))\right] \\
v_r(t+1) &= \delta \left[\gamma_I l_0(t) + \gamma_I l_1 G(V_n(t)) + w(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))}\right] \\
&\quad + f(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} + \frac{f(t)}{\phi f(t) + 1} + (1-\sigma) v_r(t)
\end{aligned}$$

Second, we assume that a falls permanently and that γ_I declines gradually after 10 years. This is the same as setting $\delta = 0$ above, shocking a, and using the following auto-regressive process for a:

$$\gamma_I = \tau_t \bar{\gamma}_I \text{ where } \tau_t = \alpha \tau_{t-1} \text{ and } \tau_1 = 1$$

We assume $\alpha = 1$ for five years when it falls to $\alpha = .95$.

These laws of motion simply add inflows and subtract outflows. A fraction $(1 - G(V_n(t)))$ of individuals who receive taste shocks default and the same fraction of individuals with taste shocks become locked in. A fraction δ of individuals who would become buyers and sellers become a buyer and a renter instead. A mass $\frac{f(t)}{\phi f(t)+1}$ experiences a foreclosure completion. The final added complexity is accounting for the mass of individuals who were locked in in period $t-1$ but are no longer locked in in period t or who were in foreclosure in period $t-1$ but are no longer in foreclosure in period t due to rising prices. Because only individuals with a loan balance above $V_n(t-1)$ are locked in at time $t-1$, this mass is a fraction $\frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))}$ of the mass $w(t)$ and $f(t)$, respectively.

These laws of motion replace the laws of motion in appendix B.1.2 above. The rest of the equations are the same, yielding a 15 equation and 15 unknown dynamic system.

B.2 Details Omitted From Main Text

B.2.1 Details of Isolating Each Effect

As described in the main text we perform three experiments to isolate the role of each driving force in our model. We provide the details of these experiments here:

1. To shut down the market tightness effect, we assume that a homeowner who defaults is not forced to become a renter for a certain random amount of time, but can instead immediately re-enter the housing market as a buyer. The law of motion for the stock of buyers in the market then becomes

$$v_b(t+1) = \gamma\left(l_0\left(t\right) + l_1\left(t\right)\right) + v_b(t)\left[1 - q_b\left(\mu\right)\sum r_m\left(t\right)e^{-\lambda(h_m(t)-a)}\right],$$

   while all other equations remain unchanged. Note that the stock of renters will always be zero in this experiment and $\mu\left(t\right) = 1$ for all t.

2. To shut down the choosey buyer effect, we assume that the buyer believes every seller he meets will be a retail seller. That is, even though there may well be distressed sellers in the market, the buyer fails to take their presence into account when determining his optimal market behavior. Along these lines, for this experiment we modify the Bellman equation of the buyer's value function to read:

$$B\left(t\right) = \beta B\left(t+1\right) + b + q_b\left(\mu\left(t\right)\right)\frac{1-\theta}{\lambda\left[1-\beta\left(1-\gamma\right)\right]}e^{-\lambda(h_n(t)-a)}.$$

   Again, we leave all other equations unchanged.

3. Finally, we run an experiment in which we include only compositional effects. For this experiment, we shut down both the market tightness and choosey buyer effects by modifying the law of motion for the stock of buyers and the Bellman equation for the buyer's value function in the manner described above.

B.2.2 Cross-Markets Analysis With 10% REO Discount

Figure 46 shows the results of the same calibration procedure in 2.7 for a 10% REO discount instead of a 20% REO discount. The lower REO discount weakens the compositional effect whereby a large REO share reduces the aggregate price index by mechanically placing more weight on properties that sell at a discount. The lower discount also weakens the choosey buyer effect, since the benefit of waiting for a foreclosure is reduced somewhat (though it still grows substantially in the downturn). Consequently, to match the non-linearity in price declines relative to the size of the preceding boom, the non-compositional

Figure 46: Cross-MSA Simulations vs. Data: 10% REO Discount

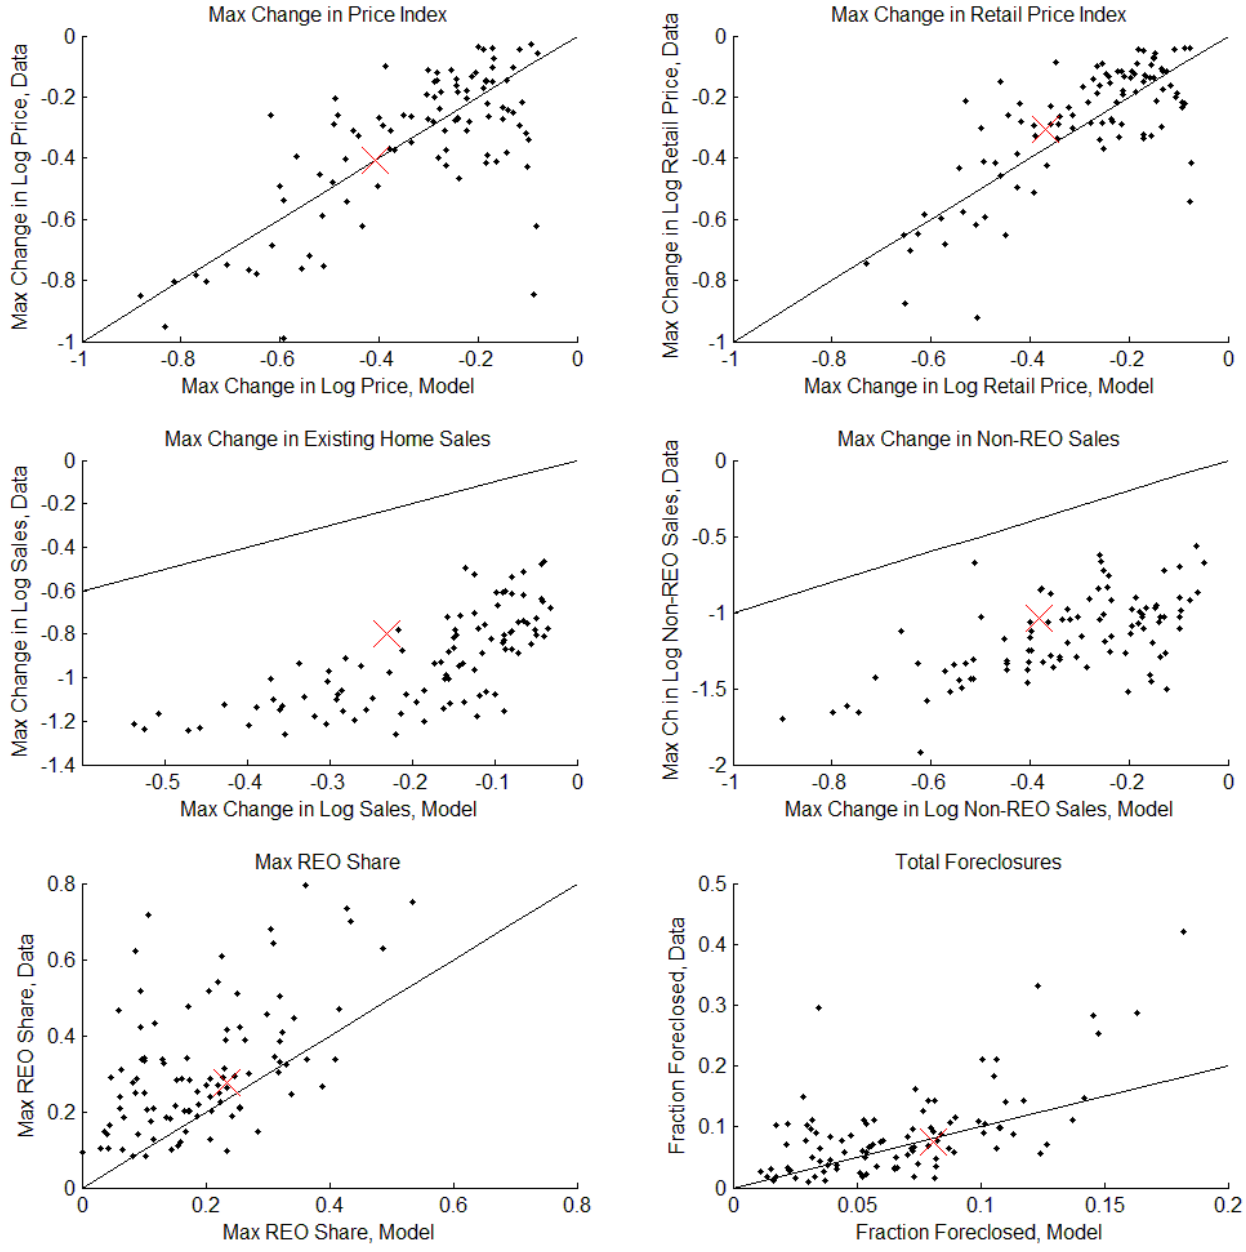

Note: Scatter plots of data vs. simulation results for 97 MSAs in regression analysis for a 10 percent discount. The red X represents the national simulation and each black dot is an MSA. The 45-degree line illustrates a perfect match between the model and the data. The variable being plotted shown in each plot's title. Data is fully described in appendix B.4. The calibration methodology described in text and appendix B.4.

Table 37: Judicial vs. Non-Judicial States

Dependent Variable	$\Delta \log(P)$	$\Delta \log(P_{Retail})$
Data	0.084	0.053
	(0.035)**	(0.025)**
Model With Backlogs	0.014	0.002
	(0.005)***	(0.002)
Model No Backlogs	0.008	0.004
	(0.004)*	(0.002)*
N	45	45

Notes: * = 10% Significance, ** = 5% Significance *** = 1% significance. All standard errors are robust to heteroskedasticity. Every reported coefficient is for the judicial state dummy in a regression that includes a linear and quadratic term for change in log price 2003-2006, z score for share with LTV over 80 percent and its interaction with change in log price 2003-2006, and z score for share with second mortgage and its interaction with change in log price 2003-2006. These regressions do not include the Saiz (2010) variables, which are not available at the state level. The columns differ by dependent variables. The rows differ by data source: the first row shows the actual CoreLogic data, the second row uses simulated dependent variable data from a model in which judicial states have a backlog, and the third row uses simulated dependent variable data from a model in which judicial states have no backlog. Every regression has 45 states as described in Appendix B.4. We use data from Mian et al. (2014), which they obtained from RealtyTrac.com, to categorize states as judicial foreclosure or non-judicial foreclosure states. Using this methodology, Connecticut, Delaware, Florida, Illinois, Indiana, Kansas, Kentucky, Louisiana, Maine, Maryland, Massachusetts, Nebraska, New Jersey, New York, North Dakota, Ohio, Pennsylvania, South Carolina, and Vermont are judicial states.

effects of foreclosure – the market tightness effect – must be larger. This results in a longer average time out of the market of 1.3 years relative to 1.05 years for a 20% steady state discount. Additionally, the permanent price decline is 22.4%, which is slightly bigger than the 21.5% for a 20% discount. Foreclosures thus exacerbate the downturn by 50%. The permanent price decline in the model remains high in order to fit the non-linearity shown in Figure 15.

Interestingly, because the compositional effect is reduced, the extent to which foreclosures exacerbate the retail price decline is increased from 28.7% to 37.5%. Intuitively, with the compositional effect weakened the retail price declines must be stronger. Unfortunately, this results in a worse average fit for retail price declines, as shown in panel B of Figure 46. Thus while the model fit for overall price declines is roughly comparable to the 20% steady state price decline case, we prefer the 20% case shown in the main text.

B.3 Judicial vs. Non-Judicial States

Table 37 shows the coefficient on judicial state of running a regression similar to equation (27) with a judicial state dummy as described in the table note.[125] The first row shows the actual data, while the second row

[125] As with the main text, the results are largely unchanged if we used weighted least squares and weight by the owner-occupied housing stock.

shows the results of a model with a backlog of $\phi = 3,400$ for judicial states and no backlog for non-judicial states and the third row shows the results of a model with no backlog for judicial or non-judicial states. Adding backlogs to the model is a step in the right direction, but the model is still an order of magnitude short of the data.

As a result, we speculate that this implies that backlogs cannot be the whole story in judicial states – there must be some reduction in the incidence of foreclosures as banks respond to the long foreclosure timelines. While it is possible that $\phi > 3,400$ in judicial states – something we cannot simulate because of numerical issues – the results of higher backlogs for a national calibration shown in Figure 23 suggests that even with a much narrower foreclosure pipeline it is not possible to get that judicial states have a log price decline that is .084 smaller from foreclosure backlogs alone.

B.4 Data Sources and Calculations

Data

The main data source is proprietary data from CoreLogic, which we supplement with data from the U.S. Census, Saiz (2010), and Mian et al. (2014).

CoreLogic provides us with a monthly data set for the nation, 50 states, and the 100 largest MSAs[126]

[126] By CBSA code and name, they are: 10420 Akron, OH; 10580 Albany-Schenectady-Troy, NY; 10740 Albuquerque, NM; 10900 Allentown-Bethlehem-Easton, PA-NJ; 12060 Atlanta-Sandy Springs-Marietta, GA; 12420 Austin-Round Rock-San Marcos, TX; 12540 Bakersfield-Delano, CA; 12580 Baltimore-Towson, MD; 12940 Baton Rouge, LA; 13644 Bethesda-Rockville-Frederick, MD; 13820 Birmingham-Hoover, AL; 14484 Boston-Quincy, MA; 14860 Bridgeport-Stamford-Norwalk, CT; 15380 Buffalo-Niagara Falls, NY; 15764 Cambridge-Newton-Framingham, MA; 15804 Camden, NJ; 16700 Charleston-North Charleston-Summerville, SC; 16740 Charlotte-Gastonia-Rock Hill, NC-SC; 16974 Chicago-Joliet-Naperville, IL; 17140 Cincinnati-Middletown, OH-KY-IN; 17460 Cleveland-Elyria-Mentor, OH; 17820 Colorado Springs, CO; 17900 Columbia, SC; 18140 Columbus, OH; 19124 Dallas-Plano-Irving, TX; 19380 Dayton, OH; 19740 Denver-Aurora-Broomfield, CO; 19804 Detroit-Livonia-Dearborn, MI; 20764 Edison-New Brunswick, NJ; 21340 El Paso, TX; 22744 Fort Lauderdale-Pompano; Beach-Deerfield Beach, FL; 23104 Fort Worth-Arlington, TX; 23420 Fresno, CA; 23844 Gary, IN; 24340 Grand Rapids-Wyoming, MI; 24660 Greensboro-High Point, NC; 24860 Greenville-Mauldin-Easley, SC; 25540 Hartford-West Hartford-East Hartford, CT; 26180 Honolulu, HI; 26420 Houston-Sugar Land-Baytown, TX; 26900 Indianapolis-Carmel, IN; 27260 Jacksonville, FL; 28140 Kansas City, MO-KS; 28940 Knoxville, TN; 29404 Lake County-Kenosha County, IL-WI; 29820 Las Vegas-Paradise, NV; 30780 Little Rock-North Little Rock-Conway, AR; 31084 Los Angeles-Long Beach-Glendale, CA; 31140 Louisville-Jefferson County, KY-IN; 32580 McAllen-Edinburg-Mission, TX; 32820 Memphis, TN-MS-AR; 33124 Miami-Miami Beach-Kendall, FL; 33340 Milwaukee-Waukesha-West Allis, WI; 33460 Minneapolis-St. Paul-Bloomington, MN-WI; 34980 Nashville-Davidson–Murfreesboro–Franklin, TN; 35004 Nassau-Suffolk, NY; 35084 Newark-Union, NJ-PA; 35300 New Haven-Milford, CT; 35380 New Orleans-Metairie-Kenner, LA; 35644 New York-White Plains-Wayne, NY-NJ; 35840 North Port-Bradenton-Sarasota, FL; 36084 Oakland-Fremont-Hayward, CA; 36420 Oklahoma City, OK; 36540 Omaha-Council Bluffs, NE-IA; 36740 Orlando-Kissimmee-Sanford, FL; 37100 Oxnard-Thousand Oaks-Ventura, CA; 37764 Peabody, MA; 37964 Philadelphia, PA; 38060 Phoenix-Mesa-Glendale, AZ; 38300 Pittsburgh, PA; 38900 Portland-Vancouver-Hillsboro, OR-WA; 39100 Poughkeepsie-Newburgh-Middletown, NY; 39300 Providence-New Bedford-Fall River, RI-MA; 39580 Raleigh-Cary, NC; 40060 Richmond, VA; 40140 Riverside-San Bernardino-Ontario, CA; 40380 Rochester, NY; 40900 Sacramento–Arden-Arcade–Roseville, CA; 41180 St. Louis, MO-IL; 41620 Salt Lake City, UT; 41700 San Antonio-New Braunfels, TX; 41740 San Diego-Carlsbad-San Marcos, CA; 41884 San Francisco-San Mateo-Redwood City, CA; 41940 San Jose-Sunnyvale-Santa Clara, CA; 42044 Santa Ana-Anaheim-Irvine, CA; 42644 Seattle-Bellevue-Everett, WA; 44140 Springfield, MA; 44700 Stockton, CA; 45060 Syracuse, NY; 45104 Tacoma, WA; 45300 Tampa-St. Petersburg-Clearwater, FL; 45780 Toledo, OH; 46060 Tucson, AZ; 46140 Tulsa, OK; 47260 Virginia Beach-Norfolk-Newport News, VA-NC; 47644 Warren-Troy-Farmington Hills, MI; 47894

that for 2000-2011 includes:

- The CoreLogic home price index and non-distressed home price index estimated from public records. We refer to these as the aggregate and retail price indices. The CoreLogic non-distressed price index differs slightly from the retail price index in the model because it excludes short sales, which we count as non-REO sales.
- The number of pre-foreclosure filings and completed foreclosure auctions estimated from public records.
- Sales counts for REOs, new houses, non-REO and non-short sale resales, and short sales estimated from public records. Because short sales are not reported separately for much of the time frame represented by the data, we combine short sales and resales into a non-REO existing home sales measure which we call retail sales. We calculate existing home sales by adding REO and retail sales. We also use this data to construct the REO share of existing home volume, which we seasonally adjust.
- Estimates of 7 quantiles of the combined loan-to-value distribution for active mortgages: under 50%, 50%-60%, 60%-70%, 70%-80%, 80%-90%, 90%-100%, 100%-110%, and over 110%. These statistics are compiled by CoreLogic using public records and CoreLogic's valuation models.
- First lien originations and first lien refinancings estimated using public records.
- Over-90-day-delinquent loans, loans in foreclosure, and active loans estimated using a mortgage-level database. We use the raw counts to construct the share of active loans that are over 90 days delinquent and in foreclosure.
- The mean number of days on the market for listed homes and closed sales estimated using Multiple Listing Service data.

We seasonally adjust the raw CoreLogic house price indices, foreclosure counts, sales counts, and delinquent and in-foreclosure loan shares using the Census Bureau's X-12 ARIMA software with an additive seasonal factor. For the state and county-level sales counts, auctions counts, days on the market, and REO share, we smooth the data using a 5 month moving average (2 months prior, the current month, and 2 months post) to remove any blips in the data caused by irregular reporting at the county level.

For the calibration of the loan balance distribution and initial number of mortgages with high LTV ratios, we adjust the CoreLogic data using data from the American Community Survey as tabulated by the Census. The CoreLogic data only covers all active loans, while our model corresponds to the entire owner-occupied housing stock. Consequently, we use the ACS 3-year 2005-2007 estimates of the owner-occupied housing

Washington-Arlington-Alexandria, DC-VA-MD-WV; 48424 West Palm Beach-Boca Raton-Boynton Beach, FL; 48864 Wilmington, DE-MD-NJ; 49340 Worcester, MA.

stock and fraction of houses with a mortgage at the national, state, and county level, which we aggregate to the MSA level using MSA definitions.[127] From this data, we construct the fraction of owner-occupied housing units with a mortgage and the fraction of owner-occupied housing units with a second lien or home equity loan. We use these estimates to adjust the loan balance distribution so it represents the entire owner-occupied housing stock and in our regressions to construct the fraction of owner-occupied houses with over 80% LTV.

The LTV data is first available for March 2006, which roughly corresponds to the eve of the housing bust as the seasonally-adjusted national house price index reached is peak in March 2006. To approximate the size of the bubble, we average the seasonally-adjusted price index for March-May 2001, March-May 2003, March-May 2006, and March-May 2011 to calculate the change in log prices for 2001 to 2006, 2003 to 2006, and 2006 to 2011. We use these variables in our regressions and to estimate the relative size of the shock for each geographical area.

We also estimate the maximum log change in seasonally-adjusted prices, smoothed and seasonally-adjusted volume, and seasonally-adjusted time to sale as well as the maximum REO share for each geographical area. We estimate the minimum value between March 2006 and December 2011 and the maximum value between January 2002 and December 2007. We implement these restrictions so that the addition of counties to the CoreLogic data set prior to 2002 does not distort our results. We calculate the fraction of the owner-occupied housing stock that was foreclosed upon by adding up completed foreclosure auctions between March 2006 and December 2012 and dividing by the owner-occupied housing stock in 2006 as calculated from the ACS adjusted for CoreLogic's approximately 85% coverage, which is assumed to be constant across locations. Again, our results are not sensitive to the choice of dates.

From the 100 MSAs and 50 states, we drop two MSAs and five states. The Birmingham, Alabama MSA is dropped because a major county stopped reporting to CoreLogic in the middle of the downturn, and the Syracuse New York MSA is dropped because loan balance distribution data is not available for this MSA in 2006. Maine, Vermont, and South Dakota are dropped because loan balance distribution data is not available for these states in 2006. For the cross-state analysis, we focus on the continental U.S. and omit Alaska and Hawaii.

We finally merge data from Saiz (2010) into the MSA data. The Saiz data includes his estimate of unusable land due to terrain, the housing supply elasticity, and the Wharton Land-Use Regulation Survey score for each MSA. We are able to match every MSA we have data on except for Sacramento and Honolulu.

Loan Balance Distribution Calibration

[127] The 3-year ACS estimates include estimates of the housing stock and houses with a mortgage for all counties with over 20,000 residents. For a few MSAs, one or more small counties are not included in the ACS data. The bias on our constructed estimates of the fraction of owner-occupied homes with a mortgage and with a second lien or home equity loan due to these small missing counties is minimal.

We use a minimum-distance methodology to calibrate the loan balance distribution for each geography. From the 7 quantiles given to us by CoreLogic and the Census Data on the number of owner-occupied homes without a mortgage, we construct a CDF of 6 points: the fraction of loans with under 50% LTV, under 60% LTV, under 70% LTV, under 80% LTV, and under 100% LTV. We then construct a norm for the distance between the Beta distribution and the empirical CDF. Because the upper tail of the distribution is most critical for our amplification channel, we weight the under 50%, under 60%, and under 70% parts of the distribution by .1 and the under 80%, under 90%, and under 100% by .2. We then choose b_a and b_b, the parameters of the Beta distribution, to minimize this norm. The resulting fit is close enough that our results are robust to alternate weightings of the norm.

Sources of Calculations in the Text

All figures in the introduction are tabulated from the CoreLogic data as described above.

For the calibration of the housing market model, the median tenure for owner occupants of approximately 9 years comes from table 3-9 of the American Housing Survey reports for 1997-2005. The 20% REO discount comes form Campbell et al.'s (2011) online appendix. They report an average discount over 1987-2009 of 26%. In table A6, they estimate this by year and show that in current housing cycle it was as low as 22.6% in 2005 and as high as 35.4% in 2009. 20 percent is thus a reasonable discount.

To determine γ_I, the incidence of income shocks for houses in negative equity, we divide the seasonally adjusted number of foreclosures by the maximum seasonally adjusted number of homes in negative equity in the CoreLogic data. The mean annual incidence is $\gamma_I = 8.6\%$.

To get that interest rates decrease the hazard of default for under-water borrowers from 8.6 percent to 7.1 percent, we use data from Bajari et al. (2010) combined with standard mortgage amortization schedules. We begin by assuming that all mortgages are at 7 percent interest rates and will be refinanced to 4 percent. The average mortgage is somewhat below 7 percent, but we choose 7 percent to reflect that some ARMs reset at quite high rates and because we want to simulate the largest possible impact of a refinancing. Assuming houses are bought with 20 percent down in steady state, a 7 percent mortgage has a monthly payment of \$1,217.50 while a 4 percent mortgage has a monthly payment of \$873.67 according to standard amortization schedules and formulas. Bajari et al. report that the mean loan in their data set has a payment-to-income ratio of .312. Assuming the \$1,217.50 monthly payment matches this ratio, monthly income is \$3,902.24. A reduction of the monthly payment to \$873.67 reduces the payment-to-income ratio by .088. Bajari et al. estimate that a one standard deviation change in the payment-to-income ratio – equivalent to a .124 change – reduces the hazard of default by 17.5 percent. Assuming linearity a reduction of .088 will reduce the hazard of default by 12.435 percent. Given the initial hazard of 8.6 percent, this implies a default hazard with a reduced interest rate of 7.1 percent.

For the principal reductions, we assume a $100 billion dollar principal reduction. 21.5 million households potentially under water implies an approximately $5,000 principal reduction for each house. If all 50 million households with a mortgage received the reduction, this would be only a $2,000 principal reduction for each house.

C Chapter 3 Appendix

C.1 Simulations of Quasi-Experiments (Figure 24)

This appendix describes the simulations of three quasi-experiments in the Rogerson and Wallenius (2009) model and the robustness of the simulations to alternative assumptions about the intensive margin labor supply elasticity. Appendix C.4 describes the analytic solution method in detail.

Calibration. The target values used to calibrate the model's parameters $\{\alpha, e_1, \bar{h}, \gamma\}$ are described in the main text. In choosing the fraction of life worked (f) for the calibration, we use the frequency at which employment is measured in the data. For instance, in the EITC simulation we calculate labor force participation in a given year as whether an individual worked at all in the past year to match the annual employment observation CPS. Because of this, the fraction of life worked at any given instant (f) differs slightly from the stated target value. To calibrate $\{\alpha, e_1, \bar{h}, \gamma\}$, we set $\gamma = \frac{1}{\varepsilon_{\mathrm{INT}}}$ to match the target for the intensive Frisch elasticity. We then calibrate the remaining parameters using the model's equilibrium conditions. Finally, we manually adjust e_1 to match $w_R/w_{\max}$, following RW.

Experiment 1: Tax Holiday in Iceland. Bianchi et al.'s data is the ratio of the total number of weeks worked to the potential supply of weeks that could have been worked by all working-age individuals in a given calendar year. We define labor force participation by whether a generation works in a given week. We then average across weeks for each calendar year to get an annual measure comparable to Bianchi et al.'s data. With $\varepsilon_{INT} = .5$, $f = 79.2\%$, $h_{\max} = .45$, and $w_R/w_{\max} = 1/2$, the calibrated values are $\gamma = 2$, $\bar{h} = .384$, $\alpha = 10.106$, and $e_1 = .593$. These parameter values generate a Frisch aggregate hours elasticity of 2.085 and a Frisch labor force participation elasticity of 1.773. These and all subsequent reported Frisch elasticities are calculated by simulating a temporary, small tax change using the same methodology as the Iceland and Canada SSP simulations; see Appendix C.4 for details. The parameter values generate a compensated aggregate hours elasticity of .663, a compensated labor force participation elasticity of .577, and a compensated intensive margin hours elasticity of .144. These and all subsequent reported compensated elasticities are calculated by comparing the steady state change in response to a small tax change; see Appendix C.4 for details. After the tax change, the maximum hours worked over the life cycle are .737 and the minimum hours worked are .570.

Experiment 2: SSP Welfare Demonstration in Canada. We generate the effective tax rates for the treatment and control groups of the SSP welfare demonstration in Canada using information on the hypothetical income of the average individual in the treatment group from Lin et al. (1998). Lin et al. use a wage regression to estimate that the predicted wage of the average individual in the treatment group is \$6.24 per hour for individuals in British Columbia and \$5.53 per hour for individuals in New Brunswick. Lin et

al. then present in Table G.2 an itemized calculation of the average treatment group individual's income accounting for taxes and other transfers under the SSP subsidy and for an individual on the standard Income Assistance (IA) welfare program. This is called hypothetical income because they use the hourly wage rate and assume the individual works 30 hours per week for 52 weeks per year in both cases.

Using this calculation, in New Brunswick an individual receiving the SSP subsidy would make $\$20,184$ per year net of taxes and transfers, while an individual working and receiving IA would make $\$14,847$ per year. If the individual did not work at all and took IA, they would not realize their earnings of $\$8,627$ but would have an IA payment that is $\$6,117$ higher. This reflects the almost dollar-for-dollar reduction of welfare payments of earnings above $\$2,400$. The individual's income would have been $\$12,337$ if they had not worked. The additional income from working $1,560$ hours per year is thus $\$2,510$ for an individual on IA and $\$7,847$ for an individual receiving the SSP subsidy. This implies an hourly wage rate of \$1.61 on IA and an effective tax rate of 70.9% under IA. Under SSP, however, the hourly wage rate is \$5.03 and the effective tax rate is 9.04%.

Similarly, for an individual in British Columbia, an individual receiving the SSP subsidy would make $\$28,267$ per year net of taxes and transfers, while an individual working and receiving IA would make $\$23,078$ per year. If the individual did not work at all and took IA, they would not realize their earnings of $\$9,734$ but would have an IA payment that is $\$7,557$ higher. The individual's income would have been $\$20,901$ if they had not worked. The additional income from working $1,560$ hours per year is thus $\$2,177$ for an individual on IA and $\$7,366$ for an individual receiving the SSP subsidy. This implies an hourly wage rate of \$1.40 on IA and an effective tax rate of 77.6% under IA. Under SSP, however, the hourly wage rate is \$4.72 and the effective tax rate is 24.3%.

Averaging the British Columbia and New Brunswick results together (as roughly half the sample resides in each area), an average single parent with one child in the control group faced effective average tax rates of 74.3% when moving from no work to full-time work at the minimum wage. An average individual in the treatment group faced effective average tax rates of 16.7% for the same change.

Card and Hyslop observe employment rates at a monthly frequency for 53 months starting from the month of random assignment. To replicate this data as closely as possible, we define labor force participation by whether a generation works in a given month. Generating $w_R/w_{\max} = 1/2$ would require $e_1 < 0$. We therefore set $e_1 = 0$, generating $w_R/w_{\max} = .615$. With $\varepsilon_{INT} = .5$, $f = 23.25\%$, and $h_{\max} = .45$, the calibrated values are $\gamma = 2$, $\bar{h} = .263$, and $\alpha = 38.378$. These parameter values generate a Frisch aggregate hours elasticity of 3.294 and a participation Frisch elasticity of 3.016. The compensated aggregate hours elasticity is .765, the compensated participation elasticity is .705, and the compensated intensive margin hours elasticity is .109. After the tax change, the maximum hours worked are .746 and the minimum hours

worked are .394.

Finally, the vast majority of individuals in the SSP sample were between the ages of 16 and 46, corresponding to the first half of life in our model. Consequently, in our simulation we only consider individuals in the first half of their life, corresponding to ages 16 to 46 out of a 60-year working life from 16 to 76.

Experiment 3: Earned Income Tax Credit in the U.S. The effective tax rates for the 1994 EITC expansion come from Meyer and Rosenbaum (2000), Table 2: the gain from working for a single mother, which includes changes in wages, welfare, Medicaid, and taxes as a result of the labor supply decision, was $\$8,943$ in 1992 and $\$10,245$ in 1996. This is relative to wages of $\$18,165$, generating effective tax rates of 50.8% in 1992 to 43.6% in 1996.

Meyer (2010) observes employment rates at an annual level using CPS data. To adjust for observables and secular time trends, Meyer regresses employment rates on observables, year dummies, and year $\times$ number of children dummies and plots the coefficients on the year $\times$ number of children dummies in Figure 2. We plot the difference between the no children dummies and a weighted average of the one child and two child dummies, using the weights reported in Table 6 of Meyer (2010). We then add the difference between the dummies and raw labor force participation rates for one- and two-child mothers to arrive at the series plotted in Figure 24c.

To replicate the data as closely as possible, in the simulations we define labor force participation by whether a generation works in a given month. Because of this, we use a target value of $f = .758$ rather than $f = .791$ as in the data. With $f = .758$ at each instant, the fraction of individuals working in each year before the quasi-experiment is approximately 79.1%. Because most single mothers are under 45, in our simulation we only consider individuals in the first half of their life, corresponding to simulated ages of 16 to 46 out of a 60-year simulated working life from 16 to 76.

With $\varepsilon_{INT} = .5$, $f = 68.7\%$, $h_{\max} = .45$, and $w_R/w_{\max} = 1/2$, the calibrated values are $\gamma = 2$, $\bar{h} = .247$, $\alpha = 22.871$, and $e_1 = .574$. These parameter values generate a Frisch aggregate hours elasticity of 2.125 and a Frisch participation elasticity of 1.814. The compensated aggregate hours elasticity is .691, the compensated participation elasticity is .608, and the compensated intensive margin hours elasticity is .144. Maximum hours worked after the tax change are .460 and minimum hours worked are .370.

Robustness. We now evaluate the robustness of the results to calibrating to an intensive margin Frisch elasticity of $\varepsilon_{\text{INT}} = .25$.

For the Iceland simulation, the calibrated values are $\gamma = 4$, $\bar{h} = .509$, and $\alpha = 32.861$. These parameter values generate a Frisch aggregate hours elasticity of 1.897 and a Frisch participation elasticity of 1.738. With these parameters, labor force participation jumps 13.3%, rather than 13.5% as presented in the main text. Maximum hours worked after the tax change are .719 and minimum hours worked are .636.

For the Canada SSP simulation, the calibrated values are $\gamma = 4$, $\bar{h} = .337$, and $\alpha = 306.149$. As above we set $e_1 = 0$, which generates $w_R/w_{\max} = .611$. These parameter values generate a Frisch aggregate hours elasticity of 3.089 and a participation Frisch elasticity of 2.949. With these parameters, labor force participation jumps from 23.5% to 76.3% one year after the subsidy is introduced. After the tax change, maximum hours worked are .585 and minimum hours worked are .421.

For the U.S. EITC simulation, the calibrated values are $\gamma = 4$, $\bar{h} = .327$, $\alpha = 179.957$, and $e_1 = .581$. These parameter values generate a Frisch aggregate hours elasticity of 1.647 and a participation Frisch elasticity of 1.475. With these parameters, labor force participation jumps from 79.1% to 85.5% on impact and then rises to 85.7% over the next 4 years. Maximum hours after the tax change are .455 and minimum hours are .409.

Calibrating to a smaller intensive Frisch elasticity of $\varepsilon_{\text{INT}} = .25$ thus does not change our conclusions: the RW model over-predicts the impacts of the temporary changes in Iceland and Canada by an order of magnitude, but is closer to matching the steady-state impact of the EITC permanent tax change.

C.2 Meta-Analysis of Quasi-Experimental Estimates (Table 17)

This appendix describes how the participation elasticities and standard errors in columns 2 and 3 of Table 17 are calculated. We report standard errors based directly on the authors' estimates if available; if not, we use the delta method to calculate a standard error for the numerator of the elasticity (log employment changes) based on reported standard errors for employment effects. If information necessary for the delta method is missing, we approximate the standard error by assuming the T-statistic on the elasticity would be the same as the T-statistic on the author's estimate.

1. Juhn, Murphy, and Topel (1991): The partial elasticity is computed by taking a weighted average of the estimates in column (3) of Table 9; the weights are computed as the fraction of the population represented by each estimate using the wage percentiles listed in column (1) of Table 9. We normalize this partial elasticity by the mean of the employment rate from 1970-89 using one minus the non-employment values reported in column (3) of Table 1. Participation is defined at the weekly level (by the fraction of weeks worked in the year). For the standard error, the variance of the partial elasticity is computed as a weighted average of the variances of the estimates in column (3) of Table 9 using the T-statistics reported in the same column. We normalize this standard error using the mean of one minus the non-employment values reported in column (3) of Table 1, assuming that non-employment is measured without error.

2. Eissa and Liebman (1996): The percentage change in participation is reported in Table III, column (4) as 2.8% with a standard error of 0.9%. The participation rate of single mothers is reported in Table II, column (1) as 73% with a standard error of 0.4%. The percentage change in net earnings for the same data

source is reported by Meyer and Rosenbaum (2000), Table 2, as the financial gain from working for single mothers in 1990 ($8,458) relative to the gain from working in 1984 ($7,469). The elasticity is thus calculated as (log(0.73+0.028)-log(0.73))/(log(8458)-log(7469)). Participation is defined as positive work hours in the past calendar year. For the standard error, the delta method is used with the additional assumption that the financial gain in the denominator, for which there is no reported standard error, is measured without error.

3. Graversen (1998): Table 5, elasticity of participation rate with respect to after tax wage, average of the four reported estimates for married women and single women, bottom panel, columns (1) and (4). The author only reports standard errors on the differences-in-differences estimates in Table 4 used to calculate the elasticities in Table 5. Because complete estimates are unavailable, we approximate the standard error of each of the four reported estimates by assuming that the T-statistics on the differences-in-difference estimates are the same as the T-statistic on the elasticities. We then average the four estimates as above to get the final reported standard error.

4. Meyer and Rosenbaum (2001): On page 1092, an elasticity of 1.07 for any employment (positive work hours) during the year is reported using gross earnings of single mothers as the base level of earnings. However, the correct denominator to calculate the percentage wage increase is net earnings prior to the reform after accounting for taxes and transfers. Making the correction requires multiplying the reported elasticity by the ratio of net earnings to gross earnings prior to the reform. Meyer and Rosenbaum (2000, page 1043) report that this ratio is 7270/18165, and thus the percentage increase in the wage is actually 45% rather than the 18% assumed to calculate the elasticity reported in Meyer and Rosenbaum (2000). The corrected elasticity estimate is given by $1.07 \times 7270/18165 = 0.43$. For the standard error, we recreate the numerator used in the calculation of the 1.07 elasticity as described by the authors on page 1091. The change in participation rate comes from the estimate in row (1), column (5) of Table 4. Base participation in 1984 and its standard error are calculated using weighted average of columns 6 and 7 of the first row of Table 2 with the weights calculated from number of observations reported in the last row of column 1 and 2 in Appendix 2. An estimate of the elasticity numerator's standard error is then calculated using the delta method. Assuming that the denominator of the elasticity and the ratio of net earnings to gross earnings are measured without error, then the numerator has the same T-statistic as the calculated elasticity. The reported standard error for the elasticity is calculated by dividing the elasticity (0.43) by the calculated numerator's T-statistic.

5. Devereux (2004): Table 4, panel 2, column (1), own-wage elasticity. Participation is defined as positive work hours in the past calendar year. Standard error from same table.

6. Eissa and Hoynes (2004): Table 6, elasticity of participation with respect wages, average estimate of

married women and married men, 2nd row from bottom. Participation is defined as positive work hours in the last year. Standard errors are calculated by recreating the authors' elasticity calculation as described on page 1951 using estimates from Table 6 and using the delta method. Base participation and wage rates are calculated from Table 2, using weighted averages of the 3rd and 4th columns based on number of observations reported in the bottom row. The reported standard error is created by combining the married women and married men standard errors as above.

7. Liebman and Saez (2006): The numerator for the elasticity is computed as log(.483-.012)-log(.483) using the Change in Wife Labor Force Participation reported in row (1) and column (2) of Table 6 and the Percent of Wives with Positive Earnings (1990-1992) reported in column (3) of Table 5. The denominator for the elasticity is computed as log(1-.419)-log(1-.31) based on the change in tax rates reported on pages 10-11 for OBRA93. Participation is defined as an indicator for positive annual earnings in the past year. Standard error is constructed using the delta method assuming that the change in tax rates is measured without error. This calculation uses the standard error on Change in Wife Labor Force Participation in Table 6 and the Percent of Wives with Positive Earnings as well as the sample size from Table 5.

8. Meghir and Phillips (2010): Page 247, last paragraph, average of single and married men in-work-income elasticities, 0.27 and 0.53 respectively. For the standard errors, the authors' calculations are replicated as described on page 247 using standard errors from Table 3.1, rows (1) and (4), column (4). The standard errors are then calculated using the delta method for each of the estimates, which are then combined to create the reported standard error.

9. Blundell, Bozio and Laroque (2013): Page 38, median overall extensive elasticity. Participation is defined as positive work hours in the past calendar year. Standard error was not reported.

10. Carrington (1996): OLS estimates from Table 2. We approximate the population-constant employment elasticity as the difference between the employment elasticity in column (1) and the population elasticity in column (5). The standard error is calculated from corresponding standard errors on elasticities under the assumption that the population and employment elasticity estimates are uncorrelated.

11. Gruber and Wise (1999): Using data reported in Table 1, the elasticity estimate is based on a regression of log(labor force participation at age 59) on log(effective net-of-tax rate) across countries. Labor force participation is defined as 1 minus fraction of Men Out of Labor Force at age 59; effective net-of-tax rate is defined as 1-implicit tax on earnings. The Netherlands is omitted from the regression because it has an implicit tax above 1. Reported standard error is from the same regression.

12. Bianchi, Gudmundsson, and Zoega (2001): Estimate and standard error from average of the elasticities for men and women reported in the text, paragraph 4, page 1570. Participation is defined at the weekly level (fraction of weeks worked in the past year).

13. Card and Hyslop (2005): From Figure 3, labor force participation before the SSP experiment is 23.6%, and the difference between the treatment and control groups during the SSP eligibility period is 13.5%. Estimated average tax rates are computed from figures in Lin et al. (1998) as described in Appendix C.1. Participation is defined as any employment in the past month. To compute standard errors, sample sizes in Table 2 adjusted for sample attrition as described in footnote 18 were combined with the data on participation rates from Figure 3. The delta method was then used assuming the change in net-of-tax wage rates was measured without error.

14. Brown (2009): We obtain an estimate of 0.08 for the elasticity of retirement age with respect to the wage using the average of the three estimates reported in column 4 of Table 2. Footnote 33 and Section 6.1 suggest that this is the author's preferred estimate. To convert this retirement age elasticity into an elasticity of years of work with respect to the wage rate, we follow footnote 30 and multiply the elasticity by the ratio of the mean age at retirement to the mean years of service reported in Table 1. The resulting elasticity is $0.08\times(60.73/26.75)$. Participation is defined as years of work, with variation on the retirement margin. The standard error is constructed from the same table and assumes that the ratio of mean age at retirement to mean years of service, for which a standard error is not reported, is measured without error.

15. Manoli and Weber (2011): Table 5, re-weighted elasticities. We first obtain separate elasticities for men and women by taking a weighted averages of the re-weighted elasticities; the weights are computed based on the fraction of individuals at each tenure threshold. The elasticity for men is 0.12 and the elasticity for women is 0.38. We then take an unweighted average of these numbers to obtain the overall elasticity of 0.25. The standard error is constructed from the same table using the same weighted average methodology.

C.3 Micro vs. Macro Elasticities (Table 18)

This appendix describes how each of the values in Table 18 are calculated. With the exception of the Frisch aggregate hours macro elasticity, the aggregate hours elasticities are defined as the sums of the intensive and extensive margin elasticities.

Hicksian, extensive margin: The micro estimate is the mean of the estimates in Panel A of Table 17. The macro estimate is computed by taking the mean of 0.13 from Davis and Henrekson (2005), 0.14 from Nickell (2003), and 0.25 from Prescott (2004). The elasticity from Davis and Henrekson is computed using the log difference in employment based on the slope coefficient in Table 3 (bottom panel, Sample C) and the sample means of labor force participation and tax rates in Table 1 for the corresponding sample. The elasticity from Nickell is computed using the average point estimate of 2 percent (reported on page 8) and the sample means of employment rates and tax rates from Tables 1 and 2, respectively. The elasticity from Prescott is calculated by regressing log labor force participation rates from OECD Stat Extracts on log net-of-tax rates

using the same sample of countries and years as Prescott.[128] The data on tax rates is taken from Table 2 of Prescott (2004). The data on labor force participation rates are missing for Canada and the U.K. in the 1970s and these observations are therefore excluded.

Hicksian, intensive margin: The micro estimate is the preferred minimum-δ estimate using Panels A and B in Table 1 of Chetty (2012). The macro estimate is the mean of the values reported by Davis and Henrekson (2005) and Prescott (2004). The value from Davis and Henrekson (2005) is computed using log differences in annual hours per employed adult based on the slope coefficient in Table 2.3 (middle panel, Sample C) and the sample means of annual hours per employed person and tax rates in Table 2.1 for the corresponding sample. The elasticity estimate can be interpreted as a compensated labor supply elasticity if government expenditure is viewed as unearned income in the aggregate. The value from Prescott (2004) is calculated by regressing log hours per worker on log net-of-tax rates using OECD data reported by Prescott in Table 2 on hours per adult, which are converted to hours per worker using labor force participation rates from the OECD Stat Extracts described above. The data on labor force participation rates are missing for Canada and the U.K. in the 1970s and these observations are therefore excluded. The elasticity estimate can be interpreted as a compensated labor supply elasticity if government expenditure is viewed as unearned income in the aggregate.

Frisch, intensive margin elasticities: the micro estimate is the unweighted mean of 0.70 in Table 2 from Pistaferri (2003) and 0.37 from Bianchi et al. (2001), as reported in Chetty (2012). The macro value in brackets is set equal to the micro estimate.

Frisch, extensive margin: The micro estimate is the mean of the estimates in Panel B of Table 1. The macro value in brackets is computed by subtracting the Frisch micro intensive margin elasticity from the Frisch aggregate hours macro elasticity.

Frisch, aggregate hours macro: the estimate is computed by taking the mean of the aggregate (total hours) elasticities implied by two models of business cycles: (1) Cho and Cooley (1994): 2.61 from the sum of the intensive and extensive margin elasticities implied by the parameters in Table 2 and (2) King and Rebelo (1999): 4 for representative agent RBC models, from page 975.

C.4 Technical Appendix

This technical appendix describes how we simulate the Rogerson and Wallenius model. We solve the model analytically as in RW (2007), the working paper version of RW (2009). All of our extensions follow RW's solution method (with slightly modified notation). Our results have been verified with iterative methods.

[128] The data are for men and women aged 15-64 for 1970-1974 and 1993-1996 in order to match Prescott's data. Data are available from OECD Stat Extracts at the following URL: http://stats.oecd.org/Index.aspx?DataSetCode=LFS_SEXAGE_I_R

The code for our simulations is available at http://obs.rc.fas.harvard.edu/chetty/index.html

Standard Rogerson and Wallenius Model. As described in the main text, each generation solves

$$\max_{c,h(a)} \log(c) - \alpha \int_0^1 \frac{h(a)^{1+\gamma}}{1+\gamma} da \text{ s.t. } c = (1-\tau) \int_0^1 e(a) \max\left\{h(a) - \bar{h}, 0\right\} da + T$$

where $e(a) = 1 - 2(1-e_1)\left|\frac{1}{2} - a\right|$ is a tent-shaped life-cycle productivity profile as shown in Figure 47. Similar to RW, we assume that the one unit of time corresponds to 60 years. We assume that time $t = 0$ corresponds to age 16, while time $t = 1$ corresponds to age 76. The model can be solved iteratively by backwards induction, but given RW's choice of functional forms it can be solved analytically as well. For consistency with RW (2007), we work with generic functions for the utility of consumption ($u(c)$), the disutility of labor supply ($v(h)$), and efficiency units of labor per hour worked (so $g(h) = \max\left\{h - \bar{h}, 0\right\}$ above) and plug in specific functional forms at the end. Each generation solves

$$\max_{c(a),h(a)} \int_0^1 u(c) - v(h(a)) da \text{ s.t. } c = (1-\tau) \int_0^1 e(a) g(h(a)) da + T$$

RW show that the optimal solution has two properties. First, there exists a cutoff e^* such that $h^*(a) > 0$ if $e(a) > e^*$ and $h^*(a) = 0$ if $e(a) \leq e^*$. Consequently, if $e(a)$ is tent shaped, there will be a date at which the individual enters the labor force and a date at which they exit, and if $e(a)$ is symmetric these dates will be symmetric around $a = .5$. Second, if $h^*(a)$ is optimal and $h^*(a_1) > 0$ and $h^*(a_2) > 0$ then $e(a_1) > e(a_2) \Rightarrow h^*(a_1) \geq h^*(a_2)$ so that the individual works weakly more hours when they have higher productivity. Finally, note that hourly wages are $w^h(a) = e(a) g(h(a)) / h(a)$.

Because individuals have a discrete labor market entry and retirement date, an individual works at all times on some interval $\left[A^E, A^R\right]$ where A^E is the labor market entry date and A^R is the retirement date. The problem can thus be rewritten:

$$\max_{c,h(a),A^E,A^R} u(c) - \int_{A^E}^{A^R} v(h(a)) da \text{ s.t. } c = (1-\tau) \int_{A^E}^{A^R} e(a) g(h(a)) da + T$$

In order to solve the model, RW re-order time, so that the most productive moment is at time 0 and the least productive moment is at time 1. Formally, define $\tilde{e}(\lambda)$ for $\lambda \in [0,1]$ so that for each λ, $\tilde{e}(\lambda)$ solves

$$\lambda = \int_0^1 I\left\{e(a) \geq \tilde{e}(\lambda)\right\} da$$

Then $\tilde{e}(\lambda)$ is the productivity level such that the individual has a productivity greater than $\tilde{e}(\lambda)$ for λ of

their life and is strictly decreasing by construction. The maximization problem can then be written as

$$\max_{c,h(\lambda),\lambda^*} u(c) - \int_0^{\lambda^*} v(h(a))\, d\lambda \text{ s.t. } c = (1-\tau)\int_0^{\lambda^*} \tilde{e}(\lambda)\, g(h(\lambda))\, d\lambda + T$$

because it will be assumed that $e(a)$ is symmetric around .5, if $\lambda^* < 1$, $A^E = .5 - \frac{\lambda^*}{2}$ and $A^R = .5 + \frac{\lambda^*}{2}$.

Under the parameters chosen by RW and that we use in our simulations, the constraint $h(a) < 1$ is always slack and can therefore be ignored. This permits an analytical solution to the problem. Plugging in the budget constraint and differentiating with respect to λ^* and $h(\lambda)$ leads to two first order conditions:

$$\frac{v(h(\lambda^*))}{u'\left((1-\tau)\int_0^{\lambda^*} \tilde{e}(\lambda)\, g(h(\lambda))\, d\lambda + T\right)} = (1-\tau)\,\tilde{e}(\lambda^*)\, g(h(\lambda^*)) \quad (65)$$

$$\frac{v'(h(\lambda))}{u'\left((1-\tau)\int_0^{\lambda^*} \tilde{e}(\lambda)\, g(h(\lambda))\, d\lambda + T\right)} = (1-\tau)\,\tilde{e}(\lambda)\, g'(h(\lambda)) \quad (66)$$

A balanced budget for the government implies that:

$$\tau\int_0^{\lambda^*} \tilde{e}(\lambda)\, g(h(\lambda))\, d\lambda = T$$

so the two FOCs can be rewritten as:

$$\frac{v(h(\lambda^*))}{u'\left(\int_0^{\lambda^*} \tilde{e}(\lambda)\, g(h(\lambda))\, d\lambda\right)} = (1-\tau)\,\tilde{e}(\lambda^*)\, g(h(\lambda^*)) \quad (67)$$

$$\frac{v'(h(\lambda))}{u'\left(\int_0^{\lambda^*} \tilde{e}(\lambda)\, g(h(\lambda))\, d\lambda\right)} = (1-\tau)\,\tilde{e}(\lambda)\, g'(h(\lambda)) \quad (68)$$

Note that if the individual works their whole life, $\lambda^* = 1$ and only the second FOC will hold. Additionally, the second $(h(\lambda))$ FOC implies that

$$\frac{v'(h(\lambda))}{(1-\tau)\,\tilde{e}(\lambda)\, g'(h(\lambda))} = u'\left(\int_0^{\lambda^*} \tilde{e}(\lambda)\, g(h(\lambda))\, d\lambda\right) = \text{constant } \forall\ \lambda \in [0,\lambda^*]$$

This differential equation pins down the entire hours profile once $h(0) = h^{\max}$ is known. Since λ^* fully pins down A^E and A^R, the optimum is defined by two free variables, $h^{\max}$ and λ^*, pinned down by the two FOCs. If $\lambda^* = 1$ due to a corner solution, the second FOC will pin down $h^{\max}$, the only free variable.

The two FOCs can be manipulated to simplify the equilibrium conditions for $h^{\max}$ and λ^*. First, divide the two FOCs to eliminate the integral and evaluate at $\lambda = 0$ to get:

$$\frac{v(h(\lambda^*))}{\tilde{e}(\lambda^*)\, g(h(\lambda^*))} = \frac{v'(h^{\max})}{e^{\max} g'(h^{\max})} \quad (69)$$

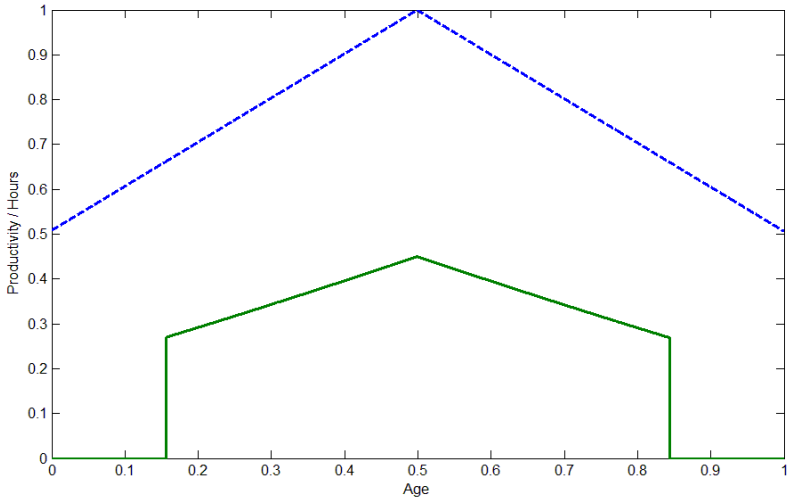

Figure 47: Productivity and Hours Profiles in the RW Model

RW show that this defines an increasing relationship between $h^{\max}$ and λ^*. Second, evaluate the second FOC at $\lambda = 0$ to get:

$$\frac{(1-\tau)\, e^{\max} g'\left(h^{\max}\right)}{v'\left(h^{\max}\right)} = \frac{1}{u'\left(\int_0^{\lambda^*} \tilde{e}\left(\lambda\right) g\left(h\left(\lambda\right)\right) d\lambda\right)} \tag{70}$$

RW show that this defines a decreasing relationship between $h^{\max}$ and λ. (69) and (70) thus together define a unique equilibrium that can be solved numerically given e_0, e_1, α, $\bar{h}$, and γ. Figure 47 illustrates the hours profile (solid green line) generated by the numerical solution using parameter values from the EITC simulation presented in the main text alongside the productivity profile (dashed blue line).

We now plug in the functional forms $u\left(c\right) = \ln\left(c\right)$, $v\left(h\right) = \alpha\frac{h^{1+\gamma}}{1+\gamma}$, $g\left(h\right) = h - \bar{h}$, and to choose a functional form for $\tilde{e}\left(\lambda\right)$. RW assume a linear formulation for the productivity profile in λ time:

$$\tilde{e}\left(\lambda\right) = \tilde{e}\left(\lambda\right) = e_0 - \left(e_0 - e_1\right)\lambda = \left(1-\lambda\right)e_0 + \lambda e_1$$

Normalizing $e_0 = 1$, this implies the an age-productivity profile of $e\left(a\right) = 1 - 2\left(1 - e_1\right)\left|\frac{1}{2} - a\right|$. With these functional forms, (69) and (70) simplify to:

$$\frac{\alpha\left(h^{\max}\right)^{\gamma}}{e_0} = \frac{\alpha h\left(\lambda^*\right)^{1+\gamma}}{\left(1+\gamma\right)\left(\left(1-\lambda^*\right)e_0 + \lambda^* e_1\right)\left(h\left(\lambda^*\right) - \bar{h}\right)} \tag{71}$$

$$\frac{\left(1-\tau\right)e_0}{\alpha\left(h^{\max}\right)^{\gamma}} = \int_0^{\lambda^*} \left(\left(1-\lambda\right)e_0 + \lambda e_1\right)\left(h\left(\lambda\right) - \bar{h}\right) d\lambda \tag{72}$$

The differential equation for hours can be manipulated to obtain

$$h(\lambda) = h^{\max}\left(\frac{(1-\lambda)e_0 + \lambda e_1}{e_0}\right)^{\frac{1}{\gamma}}$$

Plugging this into the two FOCs and simplifying gives

$$\bar{h} = \frac{\gamma}{1+\gamma} h^{\max}\left(\frac{(1-\lambda^*)e_0 + \lambda^* e_1}{e_0}\right)^{\frac{1}{\gamma}} \tag{73}$$

$$\alpha = \frac{e_0(1-\tau)}{(h^{\max})^{\gamma}\left\{h^{\max}\frac{[(1-\lambda^*)e_0+\lambda^* e_1]^{\frac{1}{\gamma}+2} - \gamma e_0^{\frac{1}{\gamma}+2}}{e_0^{\frac{1}{\gamma}}\left(\frac{1}{\gamma}+2\right)(e_1-e_0)} - \bar{h}\left[\frac{((1-\lambda^*)e_0+\lambda^* e_1)^2 - e_0^2}{2(e_1-e_0)}\right]\right\}} \tag{74}$$

The intensive margin Frisch elasticity, which is one of the moments we use for calibration, can be calculated analytically. Rearranging equation (70) and plugging in the functional forms and normalizing $e_0 = 1$ gives:

$$(1-\tau)u'(c) = \alpha(h^{\max})^{\gamma}$$

Taking logs and differentiating with respect to $1-\tau$ holding $u'(c)$ constant gives:

$$\varepsilon^{\text{Frisch}}_{h^{\max},\,1-\tau} = \frac{1}{\gamma}$$

Because the hours profile shifts vertically by $h^{\max}$ when taxes change, this is also the intensive margin Frisch elasticity in the model. Consequently, we can calibrate the model to a particular intensive margin Frisch elasticity ε_{INT} by choosing $\gamma = \frac{1}{\varepsilon_{\text{INT}}}$.

The model is calibrated as described in Appendix C.1. With $\{\alpha, e_1, \bar{h}, \gamma\}$ chosen, the model can be solved numerically by inverting equations (73) and (74) to solve for $h^{\max}$ and λ^*.

Asset Profile in the RW Model. In order to characterize the impact of unanticipated tax changes on labor supply, we need to know assets at the time of the tax change. Because assets and age are the only state variables, assets holdings are the time of the tax change are adequate to solve the model.

We assume that each generation receives a lump-sum rebate equal to the taxes they pay at each instant in time. Under this assumption, it is straightforward to back out an agent's asset position at any time. Note that the labor market entry and retirement dates are $A^E = .5 - \frac{\lambda^*}{2}$ and $A^R = .5 + \frac{\lambda^*}{2}$, respectively. Between A^E and A^R, hours are

$$h(a) = h^{\max}\left(\frac{e(a)}{e_0}\right)^{\frac{1}{\gamma}} = h^{\max}\left(\frac{e_0 - 2(e_0 - e_1)|.5 - a|}{e_0}\right)^{\frac{1}{\gamma}}$$

and so earnings when working are

$$\begin{aligned} w(a) &= g(h(a))\, e(a) \\ &= \left[h^{\max} \left(\frac{e_0 - 2(e_0 - e_1)|.5 - a|}{e_0} \right)^{\frac{1}{\gamma}} - \bar{h} \right] (e_0 - 2(e_0 - e_1)|.5 - a|) \end{aligned}$$

while consumption is always

$$c = h^{\max} \frac{[e_0 - \lambda^* (e_0 - e_1)]^{\frac{1}{\gamma}+2} - \gamma e_0^{\frac{1}{\gamma}+2}}{e_0^{\frac{1}{\gamma}} \left(\frac{1}{\gamma} + 2\right)(e_1 - e_0)} - \bar{h} e_0 \lambda^* + \bar{h}(e_0 - e_1) \frac{(\lambda^*)^2}{2}.$$

Thus assets at time t are:

$$S_t = \begin{cases} -ca, \; a < A^E \\ -ca + \int_{A^E}^{a} \left[h^{\max} \left(\frac{e_0 - 2(e_0 - e_1)|.5 - a|}{e_0} \right)^{\frac{1}{\gamma}} - \bar{h} \right] (e_0 - 2(e_0 - e_1)|.5 - a|)\, da, \; a \in \left[A^E, A^R\right] \\ -ca + c, \; a > A^R \end{cases}$$

The middle term can be simplified analytically to:

$$\begin{aligned} S_t &= -ca + h^{\max} \frac{(e_1 + 2a(e_0 - e_1))^{\frac{1}{\gamma}+2} - \left(e_1 + 2A^E(e_0 - e_1)\right)^{\frac{1}{\gamma}+2}}{2e_0^{\frac{1}{\gamma}} \left(\frac{1}{\gamma} + 2\right)(e_0 - e_1)} \\ &\quad -\bar{h} e_1 (a - A_1) - \bar{h} \left(a^2 - \left(A^E\right)^2\right)(e_0 - e_1) \end{aligned}$$

if $a \leq .5$ and

$$\begin{aligned} S_t &= -ca + S_{.5} + h^{\max} \frac{(2e_0 - e_1 - 2a(e_0 - e_1))^{\frac{1}{\gamma}+2} - (e_0)^{\frac{1}{\gamma}+2}}{2\left(2 + \frac{1}{\gamma}\right)(e_1 - e_0)} \\ &\quad -\bar{h}(2e_0 - e_1)(a - .5) + \bar{h}\left(a^2 - .5\right)(e_0 - e_1) \end{aligned}$$

if $a \geq .5$, where

$$\begin{aligned} S_{.5} &= h^{\max} \frac{(e_1 + (e_0 - e_1))^{\frac{1}{\gamma}+2} - \left(e_1 + 2A^E(e_0 - e_1)\right)^{\frac{1}{\gamma}+2}}{2e_0^{\frac{1}{\gamma}} \left(\frac{1}{\gamma} + 2\right)(e_0 - e_1)} \\ &\quad -\bar{h} e_1 \left(.5 - A^E\right) - \bar{h} \left(.5^2 - \left(A^E\right)^2\right)(e_0 - e_1) \end{aligned}$$

We solve each generation's problem separately and then add across generations, which are weighted equally, to simulate the overall response to our quasi-experiments.

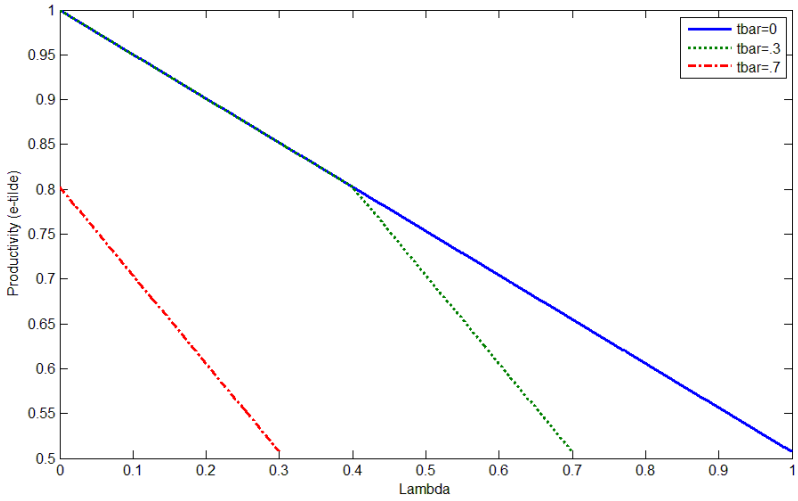

Figure 48: Productivity Profile $\tilde{e}(\lambda)$ For Various Values of Time of Tax Change $\bar{t}$

Permanent Tax Changes. We first consider the EITC simulation of a one time permanent tax change. Consider the problem of an age $\underline{t}$ individual with assets $S_{\underline{t}}$ as calculated in the previous section who faces a new tax schedule τ. The individual smooths consumption across periods and solves

$$\max_{c,h(a)} (1-\underline{t})\, u(c) - \int_{\underline{t}}^{1} v(h(a))\, da \text{ s.t. } (1-\underline{t})\, c = (1-\tau) \int_{\underline{t}}^{1} e(a)\, g(h(a))\, da + T + S_{\underline{t}}$$

This equation can be solved by analytically re-ordering time as described above in the solution to the RW model. All the solution requires is changing the $\tilde{e}(\lambda)$ profile, with $\lambda \in [0, 1-\underline{t}]$, to reflect the fact that some time has already elapsed.

The new $\tilde{e}(\lambda)$ function will be piecewise linear, as illustrated in Figure 48 using the parameter values used for the EITC simulation in the main text. When $\underline{t} = 0$, $e(\lambda) = e_0 - \lambda(e_0 - e_1)$ as above, illustrated by the solid blue line in Figure 48 below. As $\underline{t}$ rises, $e(\lambda)$ will be piecewise linear, as the low productivity time periods up to $\underline{t}$ will occur once, not twice. Thus $e(\lambda)$ will look the same for low λ, but after $2\underline{t}$ it will have twice the slope, as shown by the green dotted line in Figure 48. When $\underline{t}$ hits .5, no productivity level occurs twice and so the function will have twice the slope and be linear again. However $e^{\max}$ will fall to $e_0 - 2(e_0 - e_1)(\underline{t} - .5)$. This case is illustrated by the red dash-dot line in Figure 48.

Consequently, if $\underline{t} < .5$, $\tilde{e}(\lambda)$ is

$$\tilde{e}(\lambda) = \begin{cases} e_0 - \lambda(e_0 - e_1) & \text{if } \lambda \leq 1 - 2\underline{t} \\ 2e_0 - e_1 - 2(\underline{t} + \lambda)(e_0 - e_1) & \text{if } \lambda > 1 - 2\underline{t} \end{cases}$$

If $\underline{t} > .5$, $\tilde{e}(\lambda)$ is

$$\tilde{e}(\lambda) = 2e_0 - e_1 - 2(\underline{t} + \lambda)(e_0 - e_1)$$

With this new $\tilde{e}(\lambda)$ profile, the problem becomes

$$\max_{\lambda^* \in [0, 1-\underline{t}], h(\lambda)} (1-\underline{t}) u \left(\frac{(1-\tau)\int_0^{\lambda^*} \tilde{e}(\lambda) g(h(\lambda)) d\lambda + T + S_{\underline{t}}}{1-\underline{t}} \right) - \int_0^{\lambda^*} v(h(\lambda)) da$$

The model will have an interior solution if the tax change is not large enough to induce $h > 1$. We show this is not the case in our three applications by reporting maximum and minimum hours after the tax change for each simulation in Appendix C.1. With this constraint slack, the model can be solved analytically. Taking the first order conditions, plugging in the government's balanced budget constraint, $T = \int_0^{\lambda^*} \tilde{e}(\lambda) g(h(\lambda)) d\lambda$, and simplifying gives:

$$\frac{v(h(\lambda^*))}{u'(c)} = (1-\tau)\tilde{e}(\lambda^*) g(h(\lambda^*)) \quad (75)$$

$$\frac{v'(h(\lambda))}{u'(c)} = (1-\tau)\tilde{e}(\lambda) g'(h(\lambda)) \quad (76)$$

As in the basic RW model, the second FOC implies

$$\frac{v'(h(\lambda))}{(1-\tau)\tilde{e}(\lambda) g'(h(\lambda))} = u'(c) = \text{ constant } \forall\ \lambda \in [0, \lambda^*]$$

which pins down the hours profile.

The two FOCs can be simplified by dividing the two FOCs to eliminate the integral and evaluating at $\lambda = 0$ and by evaluating the second FOC at $\lambda = 0$. With our functional forms, this yields:

$$\frac{\alpha h(\lambda^*)^{1+\gamma}}{(1+\gamma)\left(h(\lambda^*) - \bar{h}\right)\tilde{e}(\lambda^*)} = \frac{\alpha (h^{\max})^{\gamma}}{e^{\max}}$$

$$\frac{(1-\tau) e^{\max} (1-\underline{t})}{\alpha (h^{\max})^{\gamma}} = \int_0^{\lambda^*} \tilde{e}(\lambda)\left(h(\lambda) - \bar{h}\right) d\lambda + S_{\underline{t}}$$

Finally, we know that $h(\lambda) = h^{\max}\left(\frac{\tilde{e}(\lambda)}{e^{\max}}\right)^{\frac{1}{\gamma}}$ from the differential equation for the hours profile. The two FOC simplify to:

$$\lambda^* = \tilde{e}^{-1}\left(e^{\max}\left(\frac{\bar{h}}{h^{\max}}\frac{1+\gamma}{\gamma}\right)^{\gamma}\right) \quad (77)$$

$$\frac{(1-\tau) e^{\max} (1-\underline{t})}{\alpha (h^{\max})^{\gamma}} = \frac{h^{\max}}{(e^{\max})^{\frac{1}{\gamma}}}\int_0^{\lambda^*} \tilde{e}(\lambda)^{1+\frac{1}{\gamma}} d\lambda - \bar{h}\int_0^{\lambda^*} \tilde{e}(\lambda) d\lambda + S_{\underline{t}} \quad (78)$$

which we solve numerically.

With the optimal $h^{\max}$ and λ^* in hand, it is easy to build the hours profile in calendar time. If $\lambda^* < 1-2\underline{t}$, the working life will be entirely after $\underline{t}$. The individual will enter the labor force at date $A^E = .5 - \frac{\lambda^*}{2}$ and exit at date $A^R = .5 + \frac{\lambda^*}{2}$. If $\lambda^* > 1 - 2\underline{t}$, the agent will have already started working so $A^E = \underline{t}$. They will thus exit at date $A^R = \underline{t} + \lambda^*$. To build the hours profile, we build a function $\lambda(a)$: if $\underline{t} > .5$,

$$\lambda(a) = a - \underline{t}$$

and if $\underline{t} < .5$,

$$\lambda(a) = \begin{cases} 2|a - .5| \text{ if } |a - .5| < \underline{t} \\ a - \underline{t} \text{ otherwise} \end{cases}$$

The hours profile is then generated by noting that:

$$h(a) = \begin{cases} h^{\max}\left(\frac{\tilde{e}(\lambda(a))}{e^{\max}}\right)^{\frac{1}{\gamma}}, a \in \left[A^E, A^R\right] \\ 0 \text{ otherwise} \end{cases}$$

Temporary Tax Changes. The solution method for Iceland and the Canada simulations — both of which feature a temporary tax reduction — is similar to the EITC solution. However, now there are two different periods in which the above problem is solved — one with tax τ_0 and one with tax τ_1 — and thus the solution consists of a system of four equations and four unknowns — $h^{\max}$ and λ^* in each tax regime.

Consider the problem of an age $\underline{t}$ individual with assets $S_{\underline{t}}$. From $\underline{t}$ to $\bar{t}$ they face a tax rate τ_0, and then the tax rate changes to τ_1. In this section, we assume that $\bar{t} < 1$, as if $\bar{t} \geq 1$ the individual only faces τ_0 the rest of their life and the problem reduces to the EITC problem described above. With perfect consumption smoothing, the individual's problem is:

$$\begin{aligned} & \max_{c,h(a)} (1-\underline{t})\, u(c) - \int_{\underline{t}}^{1} v(h(a))\, da \\ \text{s.t. } (1-\underline{t})\, c &= (1-\tau_0) \int_{\underline{t}}^{\bar{t}} e(a)\, g(h(a))\, da + (1-\tau_1) \int_{\bar{t}}^{1} e(a)\, g(h(a))\, da + T + S_{\underline{t}} \end{aligned}$$

Again re-order time as in RW. There will now be two $\tilde{e}(\lambda)$ functions: $\tilde{e}_0(\lambda_0)$ with $\lambda_0 \in [0, \bar{t} - \underline{t}]$ in the period with taxes τ_0 and $\tilde{e}_1(\lambda_1)$ with $\lambda_1 \in [0, 1 - \bar{t}]$ in the period with taxes τ_1. $\tilde{e}_1(\lambda_1)$ will look exactly as in the EITC simulation, with $\bar{t}$ replacing $\underline{t}$: if $\bar{t} < .5$,

$$\tilde{e}_1(\lambda_1) = \begin{cases} e_0 - \lambda_1 (e_0 - e_1) \text{ if } \lambda_1 \leq 1 - 2\bar{t} \\ 2e_0 - e_1 - 2(\bar{t} + \lambda_1)(e_0 - e_1) \text{ if } \lambda_1 > 1 - 2\bar{t} \end{cases}$$

and if $\bar{t} \geq .5$,

$$\tilde{e}_1(\lambda_1) = 2e_0 - e_1 - 2(\bar{t} + \lambda_1)(e_0 - e_1)$$

As for $\tilde{e}_0(\lambda_0)$, if $\bar{t} \leq .5$, then the area between $\underline{t}$ and $\bar{t}$ will only have the increasing side of the absolute value function:

$$\tilde{e}_0(\lambda_0) = e_0 - (1 - 2\bar{t} + 2\lambda_0)(e_0 - e_1)$$

Similarly, if $\underline{t} \geq .5$, the the area between $\underline{t}$ and $\bar{t}$ will only have the decreasing side of the absolute value function:

$$\tilde{e}_0(\lambda_0) = 2e_0 - e_1 - 2(e_0 - e_1)(\underline{t} + \lambda_0)$$

If $\underline{t} < .5$ and $\bar{t} > .5$, then we will have part of the absolute value function in the $\tilde{e}_0$. Let $t = \min\{\bar{t} - .5, .5 - \underline{t}\}$. Then

$$\tilde{e}_0(\lambda_0) = \begin{cases} e_0 - \lambda_0(e_0 - e_1) \text{ if } \lambda_0 \leq 2t \\ e_0 + 2t(e_0 - e_1) - 2\lambda_0(e_0 - e_1) \text{ if } \lambda_0 > 2t \end{cases}$$

With these profiles in hand, we note that under each tax regime an individual will always work if their productivity is above a cutoff level, as in RW. The problem can then be written as:

$$\max_{\lambda_0^* \in [0,\bar{t}-\underline{t}], \lambda_1^* \in [0,1-\bar{t}], h_0(\lambda), h_1(\lambda)} (1 - \underline{t})\, u(c) - \int_0^{\lambda_0^*} v(h_0(\lambda_0))\, d\lambda_0 - \int_0^{\lambda_1^*} v(h_1(\lambda_1))\, d\lambda_1$$

In this case, the model may not have an interior solution as an agent may find it optimal to work all of the time for which the tax is τ_0. We describe how we handle these corner solutions below.

Calculating the FOC's and plugging in the government balanced budget constraint in each period gives:

$$\begin{aligned} \frac{v(h_0(\lambda_0^*))}{u'(c)} &= (1 - \tau_0)\, \tilde{e}_0(\lambda_0^*)\, g(h_0(\lambda_0^*)) \\ \frac{v'(h_0(\lambda_0))}{u'(c)} &= (1 - \tau_0)\, \tilde{e}_0(\lambda_0)\, g'(h_0(\lambda_0)) \\ \frac{v(h_1(\lambda_1^*))}{u'(c)} &= (1 - \tau_1)\, \tilde{e}_1(\lambda_1^*)\, g(h_1(\lambda_1^*)) \\ \frac{v'(h_1(\lambda_1))}{u'(c)} &= (1 - \tau_1)\, \tilde{e}_1(\lambda_1)\, g'(h_1(\lambda_1)) \end{aligned}$$

The second FOC implies that:

$$\begin{aligned} \frac{v'(h_0(\lambda_0))}{(1 - \tau)\, \tilde{e}(\lambda_0)\, g'(h_0(\lambda_0))} &= u'(c) \\ &= \text{constant } \forall\ \lambda \in [0, \lambda_0^*] \end{aligned}$$

As before once we know $h_0(0) = h_0^{\max}$ all of $h_0(\lambda_0)$ is pinned down. Similarly, the fourth FOC implies that:

$$\frac{v'(h_1(\lambda_1))}{(1-\tau)\tilde{e}(\lambda_1)g'(h_1(\lambda_1))} = u'(c) = \text{ constant } \forall\ \lambda \in [0, \lambda_1^*]$$

We can then follow the same steps as above, dividing the two FOCs and evaluating at $\lambda_0 = 0$ and $\lambda_1 = 0$ and evaluating the second and fourth FOCs at $\lambda_0 = 0$ and $\lambda_1 = 0$, respectively. Plugging in the functional forms one gets four equilibrium conditions:

$$h_0^{\max} = \bar{h}\frac{1+\gamma}{\gamma}\left(\frac{e_0^{\max}}{\tilde{e}_0(\lambda_0^*)}\right)^{\frac{1}{\gamma}} \tag{79}$$

$$h_1^{\max} = \bar{h}\frac{1+\gamma}{\gamma}\left(\frac{e_1^{\max}}{\tilde{e}_1(\lambda_1^*)}\right)^{\frac{1}{\gamma}} \tag{80}$$

$$\begin{aligned}\frac{(1-\tau_0)e_0^{\max}(1-\underline{t})}{\alpha(h_0^{\max})^{\gamma}} = {} & \frac{h_0^{\max}}{(e_0^{\max})^{\frac{1}{\gamma}}}\int_0^{\lambda_0^*}\tilde{e}_0(\lambda_0)^{1+\frac{1}{\gamma}}d\lambda_0 - \bar{h}\int_0^{\lambda_0^*}\tilde{e}(\lambda_0)d\lambda_0 + \\ & \frac{h_1^{\max}}{(e_1^{\max})^{\frac{1}{\gamma}}}\int_0^{\lambda_1^*}\tilde{e}_1(\lambda_1)^{1+\frac{1}{\gamma}}d\lambda_1 - \bar{h}\int_0^{\lambda_1^*}\tilde{e}(\lambda_1)d\lambda_1 + S_{\underline{t}}\end{aligned} \tag{81}$$

$$\begin{aligned}\frac{(1-\tau_1)e_1^{\max}(1-\underline{t})}{\alpha(h_1^{\max})^{\gamma}} = {} & \frac{h_0^{\max}}{(e_0^{\max})^{\frac{1}{\gamma}}}\int_0^{\lambda_0^*}\tilde{e}_0(\lambda_0)^{1+\frac{1}{\gamma}}d\lambda_0 - \bar{h}\int_0^{\lambda_0^*}\tilde{e}(\lambda_0)d\lambda_0 + \\ & \frac{h_1^{\max}}{(e_1^{\max})^{\frac{1}{\gamma}}}\int_0^{\lambda_1^*}\tilde{e}_1(\lambda_1)^{1+\frac{1}{\gamma}}d\lambda_1 - \bar{h}\int_0^{\lambda_1^*}\tilde{e}(\lambda_1)d\lambda_1 + S_{\underline{t}}\end{aligned} \tag{82}$$

These four equations hold for interior solutions: $\lambda_0^* \in (0, \bar{t} - \underline{t})$ and $\lambda_1^* \in (0, 1-\bar{t})$. They also work at the $\lambda_0^* = 0$ and $\lambda_1^* = 0$ corner solutions because then the hours problem is trivial. At the $\lambda_0^* = \bar{t} - \underline{t}$ corner solution, λ_1^*, $h_0^{\max}$, and $h_1^{\max}$ are pinned down by the second, third, and fourth FOCs. At the $\lambda_1^* = 1 - \bar{t}$ corner solution, λ_0^*, $h_0^{\max}$, and $h_1^{\max}$ are pinned down by the first, third and fourth FOCs. If both λ_1^* and λ_0^* are at corner solutions, only the third and fourth FOCs apply. In each case, we solve the general four equation system and then proceed to the corner solution cases if λ_0^* or λ_1^* are not in the correct intervals. There may also be a corner solution for hours if $h_0(\lambda_0) > 1$ for some λ_0; this case is considered separately in a subsequent section.

Having solved for λ_0^*, λ_1^*, $h_0^{\max}$, and $h_1^{\max}$, we can then calculate retirement dates and build the hours profile. Let A_i^E be the labor market entry date and A_i^R be the labor market exit date under tax system i. If $\lambda_1^* < 1 - 2\bar{t}$, the working life will be entirely after $\bar{t}$. The individual will enter the labor force at date $A_1^E = .5 - \frac{\lambda^*}{2}$ and exit at date $A_1^R = .5 + \frac{\lambda^*}{2}$. If $\lambda_1^* > 1 - 2\bar{t}$, the agent will have already started working so $A_1^E = \bar{t}$. They will thus exit at date $A_1^R = \bar{t} + \lambda^*$. As for λ_0^*, if $\lambda_0^* = 0$ the worker does not work between $\underline{t}$ and $\bar{t}$. If $\bar{t} < .5$, then $A_0^E = \bar{t} - \lambda_0^*$ and $A_0^R = \bar{t}$. If $\underline{t} > .5$, then $A_0^E = \underline{t}$ and $A_0^R = \underline{t} + \lambda_0^*$. If $\underline{t} < .5$ and $\bar{t} > .5$, there are three cases. If $.5 - \underline{t} < \bar{t} - .5$, there are two cases: if $2\lambda_0 < .5 - \underline{t}$ then $A_0^E = .5 - \frac{\lambda_0^*}{2}$ and $A_0^R = .5 + \frac{\lambda_0^*}{2}$ and otherwise $A_0^E = \underline{t}$ and $A_0^R = \underline{t} + \lambda_0$. If $.5 - \underline{t} \geq \bar{t} - .5$, there are two cases: if $2\lambda_0 < \bar{t} - .5$

then $A_0^E = . - \frac{\lambda_0^*}{2}$ and $A_0^R = .5 + \frac{\lambda_0^*}{2}$ otherwise $A_0^E = \bar{t} - \lambda_0$ and $A_0^R = \bar{t}$.

In order to build the hours profile, we proceed as in the EITC section and build a $\lambda(a)$ function. $\lambda_1(a)$ looks the same as $\lambda(a)$ in the EITC simulation with $\underline{t}$ replacing $\bar{t}$. For $\lambda_0(a)$, if $\bar{t} < .5$,

$$\lambda_0(a) = \bar{t} - a$$

if $\underline{t} > .5$,

$$\lambda_0(a) = a - \underline{t}$$

If $\underline{t} < .5$ and $\bar{t} > .5$, there are two cases: if $.5 - \underline{t} < \bar{t} - .5$,

$$\lambda_0(a) = \begin{cases} 2|a - .5| & \text{if } a < 1 - \underline{t} \\ a - \underline{t} & \text{otherwise} \end{cases}$$

and if $.5 - \underline{t} > \bar{t} - .5$

$$\lambda_0(a) = \begin{cases} 2|a - .5| & \text{if } a > 1 - \bar{t} \\ \bar{t} - a & \text{otherwise} \end{cases}$$

The hours profile can then be generated from the $\lambda_0(a)$ and $\lambda_1(a)$ functions as with a permanent tax change.

Calculating Elasticities. The elasticities reported in the text and Appendix C.1 are constructed by simulation. For all of the simulations, we compare labor supply under the pre-quasi-experimental tax regime τ to labor supply under a tax regime of $\tau - .01$ to approximate an infinitesimal tax change. Denoting hours under the two tax regimes by h_1 and h_2, respectively, the elasticity is calculated as:

$$\varepsilon = \frac{\ln\left(\frac{h_2}{h_1}\right)}{\ln\left(\frac{1-\tau+.01}{1-\tau}\right)}$$

To calculate the Frisch elasticities, we treat the tax change from τ to $\tau - .01$ as a temporary tax change lasting $\frac{1}{6{,}000}$ units of time using 6,000 generations to approximate a tax change for an infinitesimal moment. Our reported elasticities are thus an approximation to an experiment in which net-of-tax wages are raised by dw for a time period dt. We report three intertemporal substitution elasticities: the intensive margin Frisch elasticity, which we know will be $\frac{1}{\gamma}$ from the derivation above, a participation Frisch elasticity, and an aggregate hours Frisch elasticity. For the aggregate hours elasticity, h_1 and h_2 are aggregate hours. For the participation elasticity, h_1 and h_2 are labor force participation rates. For the intensive margin elasticity h_1 and h_2 are aggregate hours for generations that would have supplied labor in the period of the tax change if the tax change had not occurred.

To calculate compensated elasticities, we compare the model's steady state under a tax regime of τ and a tax regime of $\tau-.01$. Our reported elasticities are thus an approximation to an experiment in which net-of-tax wages are raised permanently by dw and agents' unearned income is reduced by a commensurate amount. We report three elasticities: the intensive margin compensated elasticity, the participation compensated elasticity, and the aggregate hours compensated elasticity, which are computed in the same manner as described in the previous paragraph.

Aggregation Over Generations. The analytical methods above are used to solve for the labor supply of a given generation. We aggregate over generations to calculate the impacts of a tax change on aggregate labor supply. To approximate a continuous time environment in which a new generation is born every instant, we use numerical simulations with a large number of generations. In particular, we project the analytical solution onto a discrete-time grid for each generation, with one generation born every time period. For the Iceland simulation, we use 9,360 generations, so three generations are born or die each week. For the Canada SSP simulation, we use 7,200 generations, so 10 generations are born or die each month. For the EITC simulation, we use 6,000 generations, so 100 generations are born or die each year. We then bin the data to report the fraction of the population that worked at any point in the last week (for Iceland), month (for Canada), or year (for EITC), so that we are consistent with the quasi-experimental data. For the EITC simulation, we then aggregate up to years to reflect Bianchi et al.'s data.